

**Document Version**

Final published version

**Licence**

CC BY

**Citation (APA)**

Nguyen, T. L., de Oliveira, M. T., Braeken, A., Ding, A. Y., & Pham, Q. V. (2026). Toward Verifiable Federated Unlearning: Framework, Challenges, and The Road Ahead. *IEEE Internet Computing*, 30(1), 59-70.  
<https://doi.org/10.1109/MIC.2026.3656638>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.  
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

Article Type: Description (see below for more detail)

# Towards Verifiable Federated Unlearning: Framework, Challenges, and The Road Ahead

Thanh Linh Nguyen\*, *Trinity College Dublin, Dublin, Ireland*

Marcela Tuler de Oliveira, *Delft University of Technology, Delft, The Netherlands*

An Braeken, *Vrije Universiteit Brussel, Brussels, Belgium*

Aaron Yi Ding, *Delft University of Technology, Delft, The Netherlands*

Quoc-Viet Pham, *Trinity College Dublin, Dublin, Ireland*

*Abstract—Federated unlearning (FUL) enables removing the data influence from the model trained across distributed clients, upholding the right to be forgotten as mandated by privacy regulations. FUL facilitates a value exchange where clients gain privacy-preserving control over their data contributions, while service providers leverage decentralized computing and data freshness. However, this entire proposition is undermined because clients have no reliable way to verify that their data influence has been provably removed, as current metrics and simple notifications offer insufficient assurance. We envision unlearning verification becoming a pivotal and trust-by-design part of the FUL life-cycle development, essential for highly regulated and data-sensitive services and applications like healthcare. This article introduces VERIFUL, a reference framework for verifiable FUL that formalizes verification entities, goals, approaches, and metrics. Specifically, we consolidate existing efforts and contribute new insights, concepts, and metrics to this domain. Finally, we highlight research challenges and identify potential applications and developments for verifiable FUL and VERIFUL.*

Federated learning (FL) is a privacy-enhancing collaborative data-sharing and training paradigm in which distributed clients (e.g., end users, edge devices, enterprises, hospitals, or organizations) jointly train a global model under the coordination of a service provider (e.g., a central server/aggregator) while keeping raw data locally, thereby achieving collective intelligence [1]. FL has matured from concept to practice, empowering applications from Google keyboard next-word prediction [2] to the US cross-center cancer treatments<sup>1</sup>. Concurrently, data protection regulations, such as the EU's GDPR<sup>2</sup> and California's CCPA<sup>3</sup>, have

strengthened clients' rights to request removal of their personal data and its influence on trained models, generally called the right to be forgotten (RTBF). For example, hospitals participating in a federated diagnostic network must be able to eliminate a patient's data influence from the global model upon consent revocation. Beyond regulatory compliance, service providers also need to eliminate malicious, noisy, or unlawful data to maintain model integrity and performance. These drivers make the capability of data erasure and its associated influence a first-class requirement in FL systems.

This capability is broadly and technically formalized as machine unlearning (MUL) [3]. In MUL, while re-training from scratch offers strong completeness and guarantees that the influence of the target data<sup>4</sup> has been eliminated, it incurs prohibitive storage, com-

---

XXXX-XXX © 2025 IEEE

Digital Object Identifier 10.1109/XXX.0000.0000000

\*Part of this work was completed at TU Delft.

<sup>1</sup><https://www.canceralliance.ai/>

<sup>2</sup><https://gdpr-info.eu/>

<sup>3</sup><https://oag.ca.gov/privacy/ccpa>

---

<sup>4</sup>Target data refers to the data to be forgotten or the data being requested to unlearn by data owners/target clients.

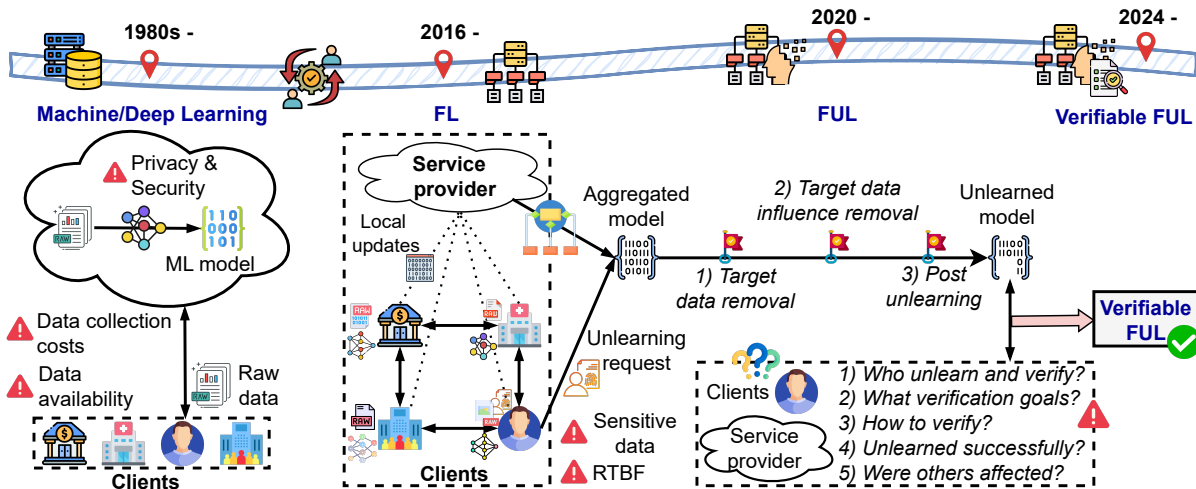


FIGURE 1: Illustration of the evolution from centralized ML to FL and FUL, highlighting key challenges at each stage and the emerging need for verifiable FUL.

putational, and time costs, especially in the era of generative AI such as large language models (LLMs). Consequently, research has shifted towards developing more efficient approximate unlearning algorithms that sacrifice the data influence removal completeness for the cost efficiency [4]. This creates an *unlearning-verification gap*, as it is difficult for service providers or clients to verify whether data influence has been removed, and as machine learning (ML) models can compress and retain knowledge from training data [5]. These challenges are further amplified by the architectural shift towards federated settings [6]. Directly adopting MUL methods is nontrivial due to FL-specific constraints, including dynamic client participation, statistical and system heterogeneity among clients, and the service provider's lack of access to raw data. While these factors have motivated research into federated unlearning (FUL) algorithms, spanning gradient modification and knowledge distillation approaches [7], underexplored questions persist from clients' and the service provider's perspectives (see Figure 1):

- 1) **Service provider & Clients:** *Who will participate in the unlearning and verification process?*
- 2) **Service provider & Clients:** *What are verification goals needed to be achieved?*
- 3) **Service provider & Clients:** *How does target data be unlearned?* (accomplished through unlearning algorithms)
- 4) **Target clients:** *How to verify that my data has been unlearned from the trained model?*
- 5) **Service provider & Clients:** *Which metrics and evidence are used to evaluate and ensure that target data is being unlearned successfully?*

- 6) **Remaining clients<sup>5</sup>:** *Will my contributions be intact?*

It is crucial to answer these questions to establish a *trust-by-design* FUL system that upholds the clients' RTBF and their right to *verify* the removal of data influence. Despite progress, a unified framework for verifiable FUL remains lacking. Romandini et al. in [7] surveyed unlearning algorithms and categorized metrics. Authors in [8], [9] outlined a workflow and a fine-grained taxonomy specifying who unlearns, who verifies, what is unlearned, and the key lessons and open directions. Gao et al. [10] proposed a mark-to-check protocol that allows the target client to verify the unlearning effect locally, using its hardest and most unique data samples. However, solely proposing unlearning algorithms and reporting metrics does not guarantee faithful unlearning execution. Critical gaps remain, including who conducts verification, what verification goals are across stakeholders (i.e., target clients, non-target clients, and service providers), how to implement verification, how well goals are achieved with selected verification approaches, and how to integrate these elements into a unified framework. If left unaddressed, target clients are often forced to unquestioningly trust the service providers' and other clients' claims about the unlearning efficacy and integrity, enabling dishonest behaviors (e.g., a service provider appears to have unlearned, passed verification, or fabricated metrics). To address this gap, we

<sup>5</sup>This article uses remaining clients and non-target clients interchangeably.

propose a unified FUL framework, called VERIFUL, which consolidates prior efforts [7]–[10] on identifying who unlearns, who verifies, and how unlearning efficacy is measured, and extends them with a systematic characterization of unlearning verification goals and verification approaches, together with metrics for assessing verification efficacy and efficiency, as well as unlearning efficacy and fidelity, within a coherent framework that we then demonstrate in a healthcare case study. Furthermore, we provide ongoing challenges and future work in this emerging domain.

## VeriFUL: VERIFIABLE FUL FRAMEWORK

This section introduces VERIFUL, specifying WHO, WHAT, HOW TO, and HOW WELL for designing a trust-by-design foundation and verifiable FUL. VERIFUL is a technical control layer that can be built on top of existing FUL systems, helping researchers and practitioners to implement established principles and obligations (e.g., GDPR's RTBF, transparency, and accountability). Figure 2 illustrates components and workflow of VERIFUL with a healthcare case study. Our focus is on the verification stage (Step 5).

### WHO - Verification Entities

Verification entities are stakeholders (e.g., clients and service providers) that participate in the verification phase to determine whether the unlearning was faithfully executed (see Step 5.1 in Figure 2).

**Clients:** comprises the target and remaining clients. Target clients submit unlearning requests and retain *RTBF* and *the right to verify* that their data<sup>6</sup> influence has been effectively removed from the global model. Remaining clients verify that their contributions (e.g., updates and data impact) to the global mode remain intact. They may act as provers (e.g., proving their updates were correctly computed) and/or as verifiers, assisting checks on target-data removal.

**Service providers:** orchestrates learning and unlearning. When executing unlearning, the service provider must generate and provide verification artifacts such as cryptographic proofs, attestations, audit logs, and evaluation metrics, demonstrating strict adherence to the unlearning algorithm. When unlearning is performed locally by target clients, the service provider may support the secure aggregation of (un-

learned) updates and assist in verifying the correctness and exclusivity of the client-side unlearning.

**Third parties:** are independent auditors (e.g., consortium authorities) serving as external validators. They inspect the artifacts supplied by the service provider and/or clients who participate in the verification process, reproduce statistical or cryptographic checks, and certify whether the stated unlearning guarantees hold.

### WHAT - Verification Goals

VERIFUL specifies guarantees and expectations a FUL system should meet, including completeness, timeliness, correctness, exclusivity, and reversibility. Clearly defining these goals establishes the foundation for systematic and reliable verification (see Step 5.2).

**Completeness:** ensures that the unlearning process eliminates target data and residual influence from both remaining clients' local models and the final global model. After unlearning, the performance on retained data of the unlearned model should be statistically indistinguishable from that of a retrain-from-scratch model trained on the same retained data, thereby preserving model utility while guaranteeing RTBF. It is noteworthy that VERIFUL mainly focuses on approximate unlearning because exact unlearning via full retraining is intractable in time and costs for modern large-scale models (e.g., LLMs) and in scenarios involving only a small number of unlearning clients or samples. Accordingly, the retrain-from-scratch model is a benchmark to quantify how closely the approximate unlearning algorithm approaches ideal completeness.

**Timeliness:** requires executing unlearning and delivering verification artifacts without undue delay following a valid erasure request. Under the GDPR, unlearning entities must act *without undue delay*, within one month of receiving the request. This deadline can be extended by up to two additional months for particularly complex cases, as long as the target clients are informed (i.e., GDPR Art. 12(3), Art. 17).

**Correctness:** guarantees that the unlearning algorithm has been carried out exactly as intended, following the prescribed protocol without deviation, omission, or adversarial manipulation. It means that the evidence provided (e.g., cryptographic proofs) must allow target clients and/or auditors to independently check unlearning protocol conformance. Unlike *completeness*, which asks whether all impacts of the target client's data were removed, correctness asks whether the removal process itself was executed faithfully and auditably.

**Exclusivity:** is the assurance that an unlearning procedure operates solely on the data associated with

<sup>6</sup>The data can be classified as personal or business ownership-related.

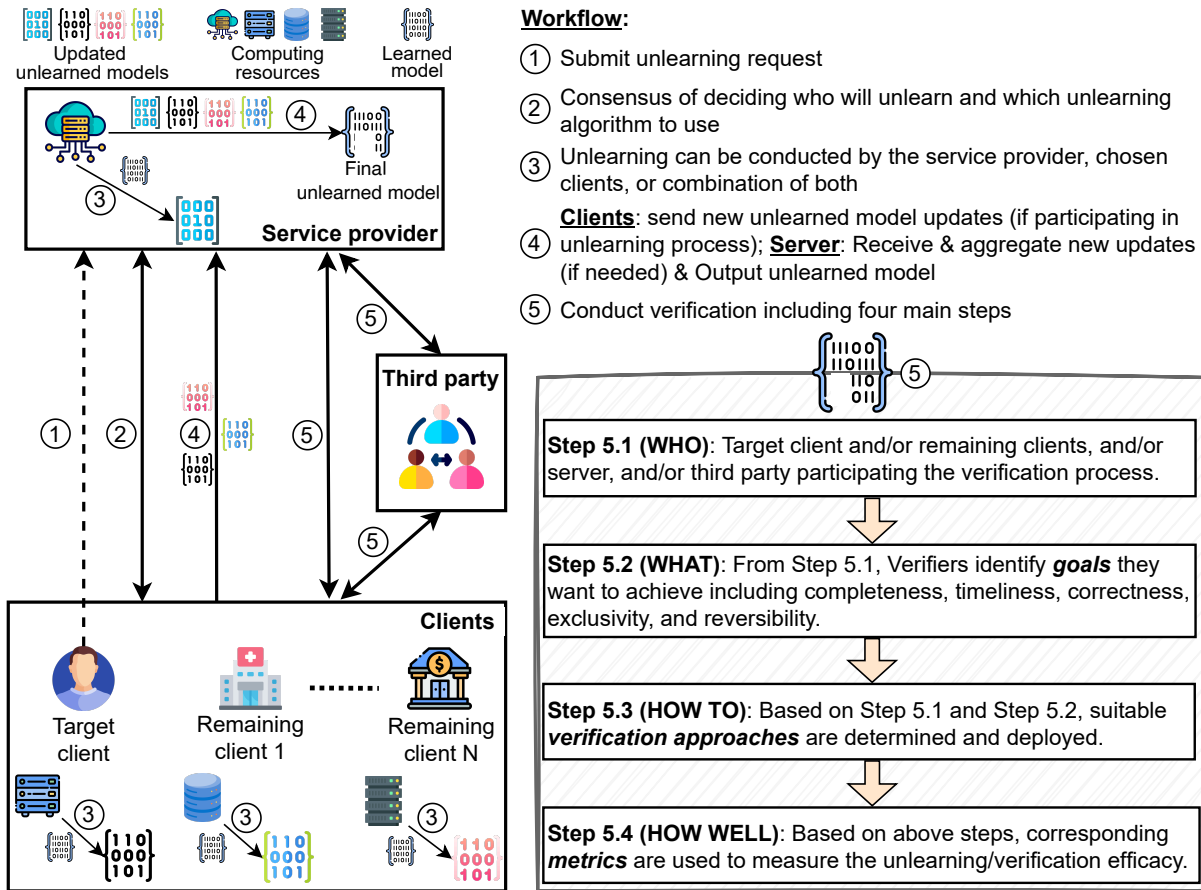


FIGURE 2: VERIFUL reference framework with workflow - we focus on Step 5 in this article.

the erasure request, preserving the integrity of all remaining clients' contributions, particularly in approximate unlearning. From the perspective of remaining clients, it addresses the fundamental concern: *Will my contributions be intact?* These concerns are reasonable because clients often have overlapping data, and their contributions also relate to their efforts' compensations (e.g., monetary rewards [11] or a shared learning model) after multiple rounds of participating.

**Reversibility:** ensures that target clients can revoke their unlearning requests, with the system efficiently restoring the forgotten knowledge in a verifiable manner, ensuring performance consistency with the pre-unlearning model.

**Discussion:** In VERIFUL, the five verification goals are not separate pillars but interdependent dimensions. Strengthening one goal may impose costs on others and must be explicitly navigated in system design. For example, tighter *completeness* coupled with stronger *correctness* usually demands deep state changes, heavier computation, log retention, and proof generation and verification, slowing down *timeliness*.

Retraining from scratch offers the strongest guarantee of eliminating all traces of target data, yet its costs and client availability constraints make it impractical in federated environments. This tension has motivated approximate unlearning methods that enhance *timeliness* while sacrificing some level of completeness [8], [12]. Likewise, mechanisms that enforce strict *exclusivity* raise a huge amount of fine-grained audits and correlated data feature and model representation checks, which require substantial verification and coordinating overhead. Enabling *reversibility* complicates the design by requiring verifiable recovery paths (e.g., weight checkpoints, client selection lists, seeds). Designing deployable systems, therefore, entails negotiating these trade-offs, balancing feasibility, forgetting efficacy, unlearning fidelity, unlearning efficiency, and the strength of guarantees. Table 1 summarizes these correlations.

## HOW TO - Verification Approaches

VERIFUL verification approaches specify methods and

TABLE 1: Qualitative Pairwise Interplay of the Five Verification Goals in VERIFUL Framework.

Goals	Completeness	Timeliness	Correctness	Exclusivity	Reversibility
Completeness	-	↓ latency (heavier compute, storage, and proof generation).	↑ correctness scheme rigor (faithful, auditable, and verifiable execution is enforced).	↓ in approximate FUL (aggressive erasure when data features among clients overlap).	↓ ease of restore (learned model checkpoints need to be recorded and retrained).
Timeliness	↓ depth forgetting (faster execution may leave residual data influence).	-	↓ guarantee strength (lightweight checks and fewer audits).	↓ non-target contribution safeguard (insufficient dependency checks).	↓ rollback fidelity (insufficient recovery paths for unlearned model).
Correctness	↑ completeness (unlearning algorithm is exactly executed without deviation).	↓ latency (proof and attestation generation and verification add overhead).	-	↑ isolation (detailed and structured logs and non-target data dependency checking).	↑ verifiable rollback (proof-based checkpoints for restoration).
Exclusivity	↓ completeness (isolation of non-target representations constrains forgetting).	↓ speed (safeguard solutions for non-target clients' contributions are integrated).	↑ correctness scheme rigor (structured logs, checkpoints, and proof-based validation).	-	↓ reversibility (restoring data target influence without perturbing others is challenging).
Reversibility	↓ completeness (a smooth restoration process limits the forgetting ability).	↓ latency (complete rollback demands longer execution time, such as proof generation).	↑ correctness (it requires fine-grained traceability and proof-based validation).	↓ exclusivity (reverting the target influence may shift decision boundaries of non-targets).	-

**Direction:** Row goal affects Column goal.

**Legend:** ↑ Improves/Strengthens, ↓ Degrades/Increases costs.

technologies that provide auditable and verifiable evidence of the unlearning procedure. This is a critical design choice as each offers distinct trade-offs in assurance strength, computational/communication overhead, scalability, and trust assumptions (see Step 5.3 in Figure 2).

**Cryptographic methods:** provide strong and mathematically verifiable proof of unlearning adherence. For instance, a zero-knowledge proof (ZKP) [13] allows a prover (e.g., the service provider or a non-target client) to convince a verifier (e.g., the target client or a trusted third party) that a statement about the unlearning process is true (e.g., the non-target clients' model updates on the non-target datasets were computed correctly using the proposed unlearning algorithm) without revealing the raw data.

Beyond ZKPs, homomorphic commitments [14] and verifiable computation systems [15] can also strengthen unlearning guarantees in federated settings. Homomorphic commitments allow for a verifiable proof that the forgotten update has been correctly subtracted from the aggregation. Similarly, verifiable computation systems such as SNARKs allow the service provider to prove that it executed the prescribed unlearning algorithm correctly when recomputing the

global parameters, so that clients can verify adherence to the protocol without re-running expensive computations.

The strength of cryptographic methods is their ability to provide strong, objective guarantees for *exclusivity* and *correctness*, which is a prerequisite for *completeness* assessments. However, a key challenge is their significant proof size, communication overhead, circuit/specification engineering, and latency impacts on *timeliness*, especially for large AI models [16].

**Hardware-based attestation:** leverages a trusted execution environment (TEE) such as Intel Software Guard Extensions [17] to perform unlearning and aggregation within an isolated enclave, providing integrity and confidentiality guarantees to the unlearning code and data even if the host system is compromised. A TEE can generate a remote attestation, a cryptographically signed report proving to a verifier that the expected unlearning code has been loaded into a genuine enclave. This contributes to *correctness* by assuring that the unlearning process is initiated with the intended code, though it does not guarantee that the code executes to completion without runtime interference or side-channel leakage. Compared to pure cryptographic proofs, TEEs achieve this with much

TABLE 2: Comparative Analysis of Verification Approaches in VERIFUL Framework

Approach	Overhead per unlearning request	Scalability w.r.t model size (M) and no. of clients (N)	Security strength	Use in VERIFUL
Cryptographic proofs	<b>High.</b> Large proof size and proof generation time.	<b>Low.</b> Degrades when M and N grow.	<b>High.</b> Strong, formal guarantees for <i>correctness</i> and <i>exclusivity</i> .	Unlearning in highly-regulated and adversarial settings with moderate M, N, and the number of unlearning requests.
Hardware-based attestation	<b>Low-moderate.</b> Near-native execution speed and costs associated with attestation.	<b>Low-moderate.</b> Limited by enclave memory limits; scales well with N but is unsuitable for large M [2].	<b>Moderate.</b> Strong guarantees for <i>correctness</i> (e.g., executing unlearning code and model updates/aggregation) inside the enclave, assuming the trusted hardware; vulnerable to side-channel attacks.	Ensure <i>correctness</i> and <i>timeliness</i> for time-critical applications with small/moderate M; execute sensitive unlearning parts.
DLT-based auditing	<b>High.</b> Significant storage and latency due to consensus and transaction history.	<b>Moderate.</b> Independent of M with off-chain storage; handle large N; ledger storage and transaction confirmation latency grow with the number of unlearning events.	<b>Moderate.</b> Strong non-repudiation and tamper-evident logging for <i>reversibility</i> ; lacks guarantees for <i>correctness</i> .	Regulated consortia requiring auditable history, accountability, and controlled rollback under moderate unlearning intensity.
Active testing	<b>Low.</b> Per-verifier overhead (i.e., unlearned model evaluation and statistical/attack analysis).	<b>High.</b> Scales well with M and N, but requires access to the final unlearned model.	<b>Low.</b> Empirical evidence towards <i>completeness</i> and <i>exclusivity</i> , but cannot guarantee <i>correctness</i> itself.	Periodic monitoring of the unlearning cycle, serving as a mandatory and lightweight complement to verification methods.

lower performance overhead, thereby supporting *timeliness*. Most current TEEs face strict enclave memory limits, which make them unsuitable for storing and processing large AI models. As a result, their use in FUL is often restricted to verifying smaller model fragments or coordinating the aggregation process.

**Distributed ledger technology-based auditing:** leverages distributed ledger technologies (DLTs), such as blockchain, to create a transparent, chronological, and immutable audit trail for the unlearning workflow. In VERIFUL, a permissioned blockchain (e.g., Hyperledger Fabric v2.x<sup>7</sup>), which offers higher throughput, lower latency, and fine-grained access control than public blockchains [18], is jointly operated by validator nodes (e.g., one or two anchor clients and a service provider), recording signed transactions that capture unlearning requests, role assignments, chosen learning/unlearning algorithms, key hyperparameters, model learning/unlearning checkpoints and their cryptographic commitments (e.g., hashes of model

updates), as well as verification outcomes or disputes. Besides, the remaining clients that perform the unlearning act as lightweight clients. These clients submit signed transactions and can verify the audit trail without participating in the heavy consensus. Rather than storing model updates on-chain, which is technically and economically costly, the blockchain only records cryptographic hashes of model updates and key artifacts, while the corresponding encrypted artifacts remain in off-chain storage [18]. During an audit or rollback, these hashes allow verifiers to check that a recovered model state is consistent with the recorded unlearning actions. Consensus mechanisms such as Practical Byzantine Fault Tolerance [19] and its variants, which are scalable with the number of unlearning requests in terms of low latency and high throughput compared to proof of work and proof of stake, and offer deterministic finality, can be an appropriate choice for DLT-based unlearning auditing. In summary, DLT-based auditing primarily strengthens *correctness* and *reversibility*, while addressing *timeliness* and integrating ZKPs for mathematical *correctness* of unlearning execution

<sup>7</sup><https://hyperledger-fabric.readthedocs.io/en/release-2.5/>

remain important directions.

**Active testing:** Distinct from DLT-based auditing, which provides immutable formal logs, active testing refers to empirical probing and assessment of the released unlearned model through evaluation and analysis. Instead of providing proof of correct execution like the approaches above, it reinforces trust by validating and providing the outcomes against targeted checks. For example, the verifier can test the unlearned model on chosen datasets (e.g., target client's data), analyze differences in model updates before and after unlearning, or repurpose security and privacy attacks as audit tools (e.g., membership inference attacks-MIAs). This approach offers flexible, intuitive, and measurable results. Solely using this approach helps us to manage *timeliness* due to its low computational overhead, but it cannot, by itself, verify *correctness*, as a malicious prover may fabricate the results without actually running the unlearning algorithm.

**Discussion:** We analyze and summarize overhead, scalability, security, and the use in VERIFUL of verification approaches in Table 2.

## HOW WELL - Verification Metrics

Verification metrics quantify how well verification goals are met under a given approach. Existing surveys summarize candidates [4], [7]–[9] but lack a unified standard for definitions, evaluation procedures, and acceptance criteria. Building on these efforts, VERIFUL consolidates, systematizes, and extends the metric space into a coherent taxonomy (see Table 3 and Step 5.4 in Figure 2).

## A Healthcare Case Study

To ground VERIFUL in a practical setting, we consider a healthcare consortium where multiple hospitals and cancer centers collaboratively train a global AI model for cancer treatment (e.g., the Cancer AI Alliance in the US<sup>8</sup>) without sharing raw patient data. An aggregation service, operated by the consortium or a service provider, coordinates training and performs model aggregation.

Suppose a participating hospital (i.e., the target hospital) needs to delete a subset of patient records and remove their influence on the global cancer treatment model. The target hospital submits an unlearning request to the consortium (Step 1). The consortium executes a consensus procedure to determine: (i) which entities will perform the unlearning (e.g., the

aggregation service, the target hospital, and/or non-target hospitals) and (ii) which unlearning algorithm will be executed (Step 2). For example, an approximate unlearning algorithm such as FUSED [20], which identifies a few knowledge-sensitive model layers, attaches sparse trainable adapters, freezes the base model, and retrains only these adapters on the retained data, can be deployed across the aggregation service, the target hospital, and non-target hospitals to produce unlearned model updates, so that retained knowledge overwrites forgotten knowledge. Unlike standard FUL, VERIFUL ensures trust by embedding verification artifacts directly into the unlearning lifecycle. The selected entities proceed to perform the unlearning procedure (Step 3), and the aggregation service then aggregates these unlearned model updates. They generate verifiable evidence, including cryptographic commitments to model updates, training data, hyperparameters, and ZKP-based proofs of correct execution.

Therefore, VERIFUL enables a robust verification stage as shown in Step 5. Here, the participating entities agree on the set of verifiers, potentially which may include the target hospital, non-target hospitals, the aggregation service, and/or independent third-party auditors (Step 5.1). These verifiers with the target hospital in agreement then define the specific verification goals (e.g., completeness, correctness), tailored to the sensitivity of the data and consortium policies (Step 5.2). Based on these goals, they utilize the pre-generated evidence to validate the unlearning process and conduct active testing (e.g., testing the unlearned cancer treatment model performance) (Step 5.3). Finally, quantitative verification metrics (e.g., performance delta of the unlearned cancer treatment model, and PVSR) are computed and analyzed to assess how well each goal has been met, thereby providing measurable assurance to the target hospital and other stakeholders regarding the forgetting efficacy, efficiency, fidelity, and integrity of the unlearning (Step 5.4). We integrate these steps and configurations as a request-level protocol around the unlearning workflow. This makes explicit how VERIFUL can guide the design of deployable verifiable FUL pipelines, while remaining agnostic to the particular unlearning algorithms and verification primitives, thereby accommodating diverse infrastructure, application requirements, and regulatory constraints.

## CHALLENGES

Given the realization of verifiable FUL remaining in its infancy, we need to tackle several challenges.

<sup>8</sup><https://www.canceralliance.ai/>

TABLE 3: Taxonomy of VERIFUL Verification Metrics.

Goal	Metric	What it measures	Typical operationalization	Acceptance criteria
Completeness	<i>Performance delta</i>	Change in learned model performance before vs. after unlearning.	Evaluate accuracy/loss/F1 score of unlearned model on the target dataset.	Significant drop (e.g., > 90%)
	<i>Residual-influence distance</i>	Divergence between unlearned model and a retrained-from-scratch model.	Compute KL/Wasserstein distance between model distributions or parameters on retained data.	Distance below predefined threshold.
	<i>Probe success rate</i>	Vulnerability to adversarial probes testing residual target data memorization.	Measures backdoor ASR on triggered patterns, or MIA accuracy and influence-function scores on target data.	Low ASR/accuracy/influence score.
Timeliness	<i>Latency</i>	Total time from request receipt to verified model availability.	Measure end-to-end time: consensus, unlearning execution, aggregation, verification (Step 2-Step 5, see Figure 2).	Meets SLA/deadline (e.g., < 24 hours and < 1 month for GDPR).
	<i>Throughput</i>	Number of unlearning requests completed per unit time.	Aggregated logs of completed requests (e.g., per second/minute/hour/day).	Throughput $\geq$ predefined unlearning target rate
Correctness	<i>Proof verification success rate (PVSR)</i>	Proportion of valid proofs accepted by verifiers.	Generate cryptographic proofs for each unlearning step; verifiers validate.	PVSR = 1.0 for all proofs. Any failure triggers audit.
	<i>Auditing score</i>	Quantitative score from third-party audit.	Auditor examines logs, checkpoints, and assigns score (e.g., 0–100).	Score above audit threshold (e.g., $\geq 90/100$ ).
Exclusivity	<i>Performance-level stability</i>	Change in performance on remaining clients' local test sets.	Evaluate unlearned model on each remaining client's fixed test set before/after unlearning.	Performance change $\leq$ small delta (e.g., $\pm 0.5\%$ accuracy).
	<i>Parameter-level stability</i>	Similarity of a remaining client's model updates before/after unlearning.	May compute cosine similarity between client's model updates.	High similarity value (e.g., $\geq 0.95$ means highly intact).
	<i>Behaviour-level stability</i>	Divergence in model output distributions on remaining client data before/after unlearning.	May compute Wasserstein distance between output distributions.	Distance value $\approx 0$ means highly intact.
Reversibility	<i>Performance consistency</i>	Performance difference between restored and pre-unlearning model.	Evaluate both models on same test set.	Change $\leq$ threshold (e.g., < 0.5% accuracy).
	<i>Restoration latency</i>	Time to revert unlearned model to its pre-unlearning state.	Measure total time from revert request to restored model availability.	Latency $\leq$ latency SLO (e.g., 30 minutes).
	<i>Model state integrity</i>	Cryptographic comparison of model states pre/post-restoration.	Compare hashes of model checkpoints (under deterministic training).	Exact hash match (if feasible).

\*\*\* KL, SLA, and SLO mean Kullback-Leibler divergence, Service Level Agreement, and Service Level Objective, respectively.

\*\*\* Setting of pre-defined thresholds, SLAs, and SLOs are dependent on application/service/existing infrastructure/client needs.

## The Co-Evolution of Unlearning Algorithms and Verifiability

The enforcement of RTBF, the need to verify, and the accountability of service providers mandate that unlearning algorithms be designed to be inherently auditable and verifiable, rather than treating verification as a post-hoc addition. This shifts the algorithm design from the sole objective of approximating the trained model to a dual goal of both achieving unlearning efficacy/efficiency/fidelity while ensuring compatibility with verification approaches. This challenge is particu-

larly for approximate unlearning. Unlike full retraining, verifying an approximation requires proving that the unlearned model operates within an acceptable/predefined threshold of residual data influence, which is ambiguous to define and audit. Consequently, the unlearning workflow must integrate detailed model transition stage logs, model checkpoints, and proof and commitment generation schemes to satisfy verification goals, such as *correctness* and *reversibility*, thereby reducing client concerns and promoting trust. However, integrating these proof-generating and model-

state management into the unlearning workflow may incur substantial latency and overhead.

### Efficiency vs. Privacy Compliance

Strict enforcement of RTBF in federated settings can conflict with the need for efficient model training and deployment. When target data is highly informative or unlearning requests occur at scale, the removal and verification processes may substantially degrade the model's generalization ability, affecting both convergence and downstream task performance. Furthermore, the additional overhead from unlearning and verification, especially using cryptographic proofs or DLTs, in terms of computation and communication, delays the release of stable global models. These challenges are further magnified in resource-constrained environments where clients have limited computing, memory, and storage to support repeated and concurrent unlearning and verifiability operations.

### Privacy and Security

In the current landscape, a client's intent to unlearn becomes explicitly disclosed through direct requests, which exposes the request to the service provider and other clients. While necessary for coordination and auditability, this transparency opens the door for adversarial behaviors. Particularly, the service provider may selectively omit, delay, or ignore unlearning requests, thereby undermining both the confidentiality and enforceability of the unlearning process. Moreover, using *active testing-based verification* (e.g., backdoor auditing and MIAs) may introduce unintended vulnerabilities [10]. Malicious verifiers (e.g., the service providers, non-target clients) could exploit these techniques to extract sensitive information, infer training set membership, or tamper with model behavior.

### Incentive mechanisms

Designing incentive mechanisms for verifiable FUL remains a significant challenge. Clients may exhibit strategic or inconsistent behaviors, causing a dynamic environment and carryover effects, such as submitting unlearning requests, revoking them later, participating selectively in verification phases, or engaging in learning only when it directly benefits them. These behaviors introduce significant inefficiencies to the service provider and other clients, including wasted computational resources, delayed model updates, and increased coordination overhead. Furthermore, the costs of unlearning and verification are often unevenly distributed among participants and private to the service

provider, creating information asymmetry among them, complicating the incentive mechanism design.

## THE ROAD AHEAD

Building on VERIFUL, we now identify four promising research directions for future investigation.

### Application-specific Solutions

A one-size-fits-all VERIFUL is infeasible. Solutions should be tailored to the specific application requirement. In highly regulated, privacy-sensitive sectors such as healthcare, an unlearning request is often driven by a client's exercise of their RTBF. The system must prioritize provable *completeness* to ensure the thorough removal of data influence and strict *correctness* to provide an auditable trail for regulatory compliance, accepting weaker *timeliness* if necessary. When unlearning is service provider-initiated to remove illegal or malicious data impacts, the primary objective is to maintain model integrity and inference performance. Verification, therefore, focuses on *completeness* to remove the harmful impact and *exclusivity* to protect remaining clients' contributions, alongside empirical metrics that validate the model's inference post-unlearning. Some applications require *reversibility* to accommodate clients' revoked requests (e.g., opt-in/opt-out on an online federated education platform with federated personalization). Ultimately, designing a deployable system requires negotiating the trade-offs between the verification goals outlined in the VERIFUL framework to align with each application's specific legal, ethical, and operational context, motivating standardized benchmarks that capture these differing objective profiles.

### Target Client-Executed Unlearning

While current verifiable FUL research typically assumes that unlearning is performed by the service provider and/or remaining clients [8], we envision an emerging direction that shifts the unlearning computation to the target clients. This client-centric approach alters the trust model for verification by reducing dependence on the service provider and other clients, while providing target clients greater control over unlearning processes. However, it introduces challenges in verifying that the target client's local unlearning computation was performed *correctly* and *exclusively* on the target data. Malicious clients may exploit this process to inject poisoned updates or tamper with the global model under the guise of unlearning, posing risks to both the service provider and other clients. Moreover, FUL

clients may be resource-constrained, which hinders their local unlearning and verification execution. Therefore, it is necessary to develop lightweight and scalable verification protocols for client-side unlearning.

### Incentive-guided Unlearning and Verification

Without fair and transparent incentives, a verifiable FUL system is likely unsustainable. Frequent unlearning requests with high intensity (e.g., removal of large data volumes driven by strict privacy preferences) may cause irreversible degradation in model performance. Future research should explore incentive mechanisms that reward participants in proportion to their verifiably measured resource contributions while ensuring system-wide fairness. Particularly, designing game-theoretic-based incentive mechanisms can enable service providers and target clients to negotiate acceptable trade-offs between privacy preferences, performance degradation, and economic rewards. Such mechanisms sustain participation in training, unlearning, and verification. For instance, a transparent incentive layer atop VERIFUL framework using blockchain. Client contributions to unlearning and verification are recorded on-chain, enabling automated auditing and reward allocation via smart contracts.

### Predictive Scaling Laws for Verifiable FUL

An open direction is an empirical framework to quantify the cumulative impact of unlearning on global model utility (e.g., performance, verification efficacy). Applying scaling-law analysis can characterize how model utility varies with unlearning intensity (e.g., volume, type, distribution of target data, and request dynamics). Then, saturation points are identified, which further forgetting leads to catastrophic utility loss or renders verification metrics unreliable. These insights enable service providers and clients to make informed trade-offs between RTBF and preservation of model utility.

## CONCLUSION

In this article, we have introduced verifiable FUL (VERIFUL), a new paradigm in which verification is integrated by design to uphold RTBF and the right to verify, thereby strengthening client trust and control of privacy. We have proposed a reference verifiable FUL framework VERIFUL, detailing the verification entities, multifaceted goals, technical verification approaches, and quantitative metrics. We have presented the key challenges to enable the practical deployment of verifiable FUL. Finally, we have highlighted promising research directions to guide future advancements in this domain.

## REFERENCES

1. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*, 2017, pp. 1273–1282.
2. A. Hard *et al.*, "Federated learning for mobile keyboard prediction," *arXiv preprint arXiv:1811.03604*, 2018.
3. Y. Cao and J. Yang, "Towards making systems forget with machine unlearning," in *IEEE Symposium on Security and Privacy*, 2015, pp. 463–480.
4. N. Li, C. Zhou, Y. Gao, H. Chen, Z. Zhang, B. Kuang, and A. Fu, "Machine unlearning: Taxonomy, metrics, applications, challenges, and prospects," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 8, pp. 13 709–13 729, 2025.
5. N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, and D. Song, "The secret sharer: Evaluating and testing unintended memorization in neural networks," in *28th USENIX Security Symposium*, 2019, pp. 267–284.
6. V.-T. Tran, H.-H. Nguyen-Le, and Q.-V. Pham, "ToFU: Transforming how federated learning systems forget user data," in *European Conference on Artificial Intelligence (ECAI)*, 2025.
7. N. Romandini, A. Mora, C. Mazzocca, R. Montanari, and P. Bellavista, "Federated unlearning: A survey on methods, design guidelines, and evaluation metrics," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 7, pp. 11 697–11 717, 2025.
8. Z. Liu *et al.*, "A survey on federated unlearning: Challenges, methods, and future directions," *ACM Computing Surveys*, vol. 57, no. 1, pp. 1–38, 2024.
9. H. Jeong, S. Ma, and A. Houmansadr, "A survey on federated unlearning: Challenges and opportunities," *arXiv preprint arXiv:2403.02437*, 2025.
10. X. Gao, X. Ma, J. Wang, Y. Sun, B. Li, S. Ji, P. Cheng, and J. Chen, "VeriFi: Towards verifiable federated unlearning," *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 6, pp. 5720–5736, 2024.
11. Q. Wang, R. Xu, S. He, R. Berry, and M. Zhang, "Unlearning incentivizes learning under privacy risk," in *Proceedings of the ACM on Web Conference 2025*, 2025, pp. 1456–1467.
12. T. T. Nguyen *et al.*, "A survey of machine unlearning," *ACM Transactions on Intelligent Systems and Technology*, Jul. 2025.
13. S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," in *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*, 2019, pp. 203–225.

14. T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *Advances in Cryptology - CRYPTO '91*, vol. 576, 1991, pp. 129–140.
15. B. Parno, C. Gentry, J. Howell, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," in *IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 238–252.
16. Z. Xing *et al.*, "Zero-knowledge proof-based verifiable decentralized machine learning in communication network: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, 2025.
17. F. McKeen *et al.*, "Innovative instructions and software model for isolated execution," in *International Workshop on Hardware and Architectural Support for Security and Privacy*, ser. HASP '13, 2013.
18. T. L. Nguyen *et al.*, "Blockchain-empowered trustworthy data sharing: Fundamentals, applications, and challenges," *ACM Computing Surveys*, vol. 57, no. 8, pp. 1–36, 2025.
19. X. Zuo, M. Wang, T. Zhu, L. Zhang, S. Yu, and W. Zhou, "Federated Learning With Blockchain-Enhanced Machine Unlearning: A trustworthy approach," *IEEE Transactions on Services Computing*, vol. 18, no. 3, pp. 1428–1444, 2025.
20. Z. Zhong, W. Bao, J. Wang, S. Zhang, J. Zhou, L. Lyu, and W. Y. B. Lim, "Unlearning through knowledge overwriting: Reversible federated unlearning via selective sparse adapter," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 30661–30670.