

Research Paper



FedSSL: privacy-preserving federated self-supervised learning with differential privacy guarantees for heterogeneous edge environments

Zayyanu Yunusa*^{id}

*Computer Science, Iconic Open University of Nigeria, Iconic Open University of Nigeria, Bakura, Nigeria.

Article Info

Article History:

Received: 25 February 2025

Revised: 05 April 2025

Accepted: 13 April 2025

Published: 02 June 2025

Keywords:

Federated Learning


Self-Supervised Learning

Differential Privacy

Contrastive Learning

Non-IID

Intelligence



ABSTRACT

In this study, we delve into the significant impact of AI, investigating its multifaceted consequences on society. In a rapidly evolving digital landscape, the integration of artificial intelligence has emerged as a transformative force in reshaping human progress and well-being. Drawing from historical perspectives, we trace the evolution of AI and its transformative journey. Our research aims to comprehensively analyze approaches for fostering responsible AI growth while mitigating potential hazards. By illuminating both the promise and perils of AI, this study contributes to informed decision-making in the unfolding AI era. This research explores the innovative utilization of AI technologies to optimize individual and societal gains while concurrently influencing the reconstruction of prevailing social norms. By harnessing the power of AI, we embark on a journey toward a new era, characterized by more efficient problem-solving, enhanced decision-making, and the redefinition of traditional social paradigms. This study investigates the multifaceted impacts of AI on human development, offering insights into its potential to revolutionize our world and foster a future marked by unprecedented advancements in progress, wellness, and social transformation.

Corresponding Author:

Zayyanu Yunusa

Computer Science, Iconic Open University of Nigeria, Iconic Open University of Nigeria, Bakura, Nigeria.

Email: yzayyanu@gmail.com

Copyright © 2025 The Author(s). This is an open access article distributed under the Creative Commons Attribution License, (<http://creativecommons.org/licenses/by/4.0/>) which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

As an example, the increasing ubiquitousness of various intelligent edge devices, such as smartphones, wearables, medical sensors and industrial IoT platforms, has created unparalleled mountains of distributed, privacy-oriented data that must adhere to the regulation, such as the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA) and the EU AI Act [1]. The current centralized machine learning (ML) approach, where raw data is aggregated on a central server, is not allowed by law in numerous geographies and not possible under the data sovereignty restrictions.

The challenge is met by Federated Learning (FL) [2] which involves a federated network of clients training a model in collaboration by only sharing model parameters or gradients, not the actual data. But there are two basic and interrelated challenges to practical FL deployments. First, in real world situations, datasets of clients are not identically distributed locally, are heterogeneous and non-stationary and in many cases non-independent (non-IID) with respect to each other, which means that when training locally, the client drifts from the original data, and when testing locally, the performance degrades when the clients are aggregated globally [3]. Second, gradient-based privacy attacks have been proven to be very effective: [4] showed that private training samples can be reconstructed with high fidelity directly from shares of gradients, thus violating the privacy assumptions of FL.

At the same time, self-supervised learning (SSL) has taken the center stage in the field of representation learning. Contrastive methods like SimCLR [5], MoCo [6] and BYOL [7] have also performed well in comparison to full data supervised methods without the use of the labeled data. For this reason, integrating SSL into federated frameworks is an appealing option since it can overcome the limitations of a lack of labels for heterogeneous clients while utilizing local data that is rich in unlabeled data. While the situation of federated SSL, non-IID heterogeneity and formal differential privacy guarantees is still significantly under-explored, this is a vital one. Currently available federated SSL methods, such as those in [8], [9] do not guarantee formal DP, and the simultaneous non-IID and DP convergence is not well understood theoretically.

Differential privacy (DP) [10] is a well-defined mathematical tool to limit the information disclosure of individual data items. DP-FL methods [11] add calibrated Gaussian or Laplace noise to the gradient updates, which however has a naive uniform noise injection that affect utility, particularly in high-dimensional gradient spaces typical in deep neural networks. In non-IID settings, where SNR is further reduced due to “client drift”, the need for privacy vs. model utility is further complicated [12].

This paper proposes a unified federated self-supervised learning framework called FedSSL that overcomes these challenges. FedSSL is contributing in the following four main ways:

- 1. FedSSL Framework:** A federated self-supervised learning model with a federated self-supervised learning framework by incorporating local contrastive NT-Xent objectives and a momentum-updated global prototype distillation mechanism for transferable representation learning across non-IID clients without labelled data.
- 2. Selective Gradient Perturbation (SGP):** A novel DP mechanism that uses an empirical Fisher Information Matrix (FIM) approximation based on the diagonal to identify the privacy-sensitive gradient subspaces, focusing Gaussian noise on top those subspaces while reducing noise in the orthogonal ones significantly enhances the privacy-utility frontier.
- 3. Formal Convergence Guarantees:** Theoretical results showing that FedSSL is $O(1/\sqrt{T})$ convergent under simultaneous non-IID data heterogeneity and (ϵ, δ) -DP constraints and with DP penalty reduced by SGP subspace selection.
- 4. Comprehensive Empirical Validation:** State-of-the-art results for five different, but similar heterogeneous benchmark domains (healthcare, image classification, activity recognition, and natural language processing) that establish its superiority over all evaluated federated privacy-aware baseline models.

The rest of this paper follows the structure below. In Section 2, the literature related to the present work is discussed. The FedSSL methodology is discussed in Section 3. The results, discussion and ablation are detailed in Section 4. The paper is concluded in section 5.

2. RELATED WORK

2.1 Federated Learning under Non-Iid Data

The baseline FedAvg algorithm [2] works well for IID-case, but suffers from serious difficulties when the client distributions are not IID. In [3] the first mathematical non-IID convergence bounds for FedAvg are introduced, which brings into focus the dependence of convergence with respect to client heterogeneity. To prevent client drift, FedProx [13] added a proximal regularization term to ensure that the changes made by the client are kept within a certain bound. SCAFFOLD [14] is an extension of the original algorithms that explicitly correct for gradient variance from partitioning that is non-IID resulting in tighter convergence guarantees. FedNova [15] and FedDyn [16] also use adaptive gradient correction mechanisms to tackle the issue of objective inconsistency due to inconsistent local updates. A few works have also been done to extend federated optimization to continual learning, where knowledge transfer is performed in a weighted manner between federated clients [17] and even [18] introduced a clustered federated learning approach to address extreme statistical heterogeneity by applying clustering on data distributions among federated clients. All these approaches are however, supervised, and not able to include SSL or formal DP.

2.2 Differentially Private Federated Learning

[10] Was the first formalizer of differential privacy, which also laid the foundations of theory of privacy-preserving computation. [11], [12] proposed DP-FedSGD & DP-FedAvg, respectively, where they added Gaussian noise to the client gradient updates before aggregation. To meet the above requirement for iterative federated training, [19] introduced Rényi Differential Privacy (RDP) which allows tighter composition bounds over multiple communication rounds. Sparse gradient methods [20] use DP noise on the top-k components of the gradient, though are not based on any principled selection of the gradient components [21]. They also further explored sparse gradient compression for distributed learning, and showed that although it decreases the communication cost, it increases the approximation error under noise. Building on this theoretical work, we find in [22] refined convergence bounds for FedProx under the local dissimilarity condition, giving a non-smoothness result that complements our theoretical analysis. This work is extended in our SGP mechanism, where we use the Fisher Information Matrix to allocate noise to the true privacy sensitive gradient subspaces as opposed to the magnitude based sparsity.

2.3 Self-Supervised and Contrastive Learning

In SimCLR [5] they introduced contrastive self-supervised learning by adding some augmented positive pairs, and it was concluded that robust data augmentation and large batch sizes lead to good representation quality. We present MoCo, a momentum encoder that decouples the batch size from the number of negative samples [6].

In [7] it is proposed that negative samples are completely avoided by bootstrapped online-target learning utilizing a momentum-updated target network. The SSL is adapted to federated settings by [8], [9] with the FedSimCLR and MOON respectively, the latter using contrastive objectives between local and global model representations. However, neither of them offers formal DP guarantee or non-IID plus DP theoretical analysis. The integration of privacy, heterogeneity, and self-supervision are acknowledged as some of the key challenging areas of federated learning [23]. To meet this need, FedSSL introduces a new approach that combines momentum distillation with privacy-preserving gradient perturbation based on SGP in a unified convergence-theoretic framework.

3. METHODOLOGY

3.1 Problem Formulation

Suppose that there are $N = 10$ “clients” with their own local data set D_i that are sampled from local distribution P_i . We have a semi-supervised problem environment, where for most clients only about 5% of the samples is labeled, and the majority is unlabeled. The objective of the global optimization is:

$$\min_{\theta} \sum_{i=1}^k p_i \cdot L_i(\theta; D_i), \text{ where } p_i = |D_i| / \sum_j |D_j|$$

Is subjected to (ϵ, δ) differential privacy constraints: No raw data or information that compromises privacy leaves from any client device. The P_i distributions are assumed to be heterogeneous (i.e., non-IID) and modeled using Dirichlet partitioning, with concentration parameter $\alpha_{Dir} = 0.1$, as suggested by the experimental protocol in [24].

3.2 Local Contrastive Objective

Data augmentation is stochastic in each client, with clients producing two correlated views (x^+, x^-) from each x . Both views are encoded in an online encoder f_{θ} and g_{ϕ} and normalized to z^+ and z^- . The local NT-Xent (normalized temperature-scaled cross-entropy) contrastive loss over mini-batch B is given by:

$$L_{con} = -\sum_i \log \left[\frac{\exp(\text{sim}(z_i^+, z_i^-)/\tau)}{\sum_{k \neq i} \exp(\text{sim}(z_i^+, z_k^-)/\tau)} \right]$$

Where $\text{sim}(\bullet, \bullet)$ is the cosine similarity and $\tau = 0.07$ is the hyperparameter of the temperature, tuned to maximize the uniformity of the representations. This objective could be capably expressed as an incentive for representations of the same sample to be close together, and representations of different samples to be farther apart.

3.3 Momentum Distillation from Global Prototypes

To offset client drift that can result from client distributions not being independent and identically distributed, FedSSL provides a global prototype alignment mechanism. To ensure global prototype embeddings $\{p_c\}$ are computed for each cluster c , using an exponential moving average (EMA) of the aggregated weights of the encoder across the rounds. Each client broadcasts prototype embeddings in each round of the communication process. For each client, it is necessary to reduce another distillation loss:

$$L_{dist} = \sum_x \| f_{\theta}^g(x) - \text{sg}(p_{\{c(x)\}}) \|_2^2$$

Here, $\text{sg}(\bullet)$ represents the stop-gradient operator preventing gradient propagation from the prototype to the local encoder, while $p_{\{c(x)\}}$ is the global prototype for the cluster $c(x)$ of input x and is picked using online k -means clustering procedure. This is the sum of the two terms:

$$L_{local} = L_{con} + \lambda \cdot L_{dist}, \text{ with } \lambda = 0.5$$

The momentum distillation term can be treated as a regularization term that pulls local representations towards prototypes that are shared across all clients, minimizing divergence between the different clients without sharing raw data.

3.4 Selective Gradient Perturbation (SGp)

Standard DP-FL randomly adds Gaussian noise to all gradient components, which inappropriately adds large amounts of noise to low information density components, while adding minimal noise to high information density components. The information density is not uniform across the gradient space, and thus it is unnecessarily wasteful to add a high amount of noise to components with low information density. Here, SGP tackles the issue by estimating the per-component privacy sensitivity through the diagonal Fisher Information Matrix (FIM).

Specifically, each client calculates an approximation to the FIM of the client with the linear map $\text{diag}(\nabla L_i \circ \nabla L_i)$ to compute. Specifically, each client computes an approximation of the FIM of the linear map $\text{diag}(\nabla L_i \circ \nabla L_i)$. The upper α fraction of gradient components with the largest Fisher sensitivity is called the subspace S , which is applied the noise σ_{high} ; the complementary subspace S^c with components having smaller Fisher sensitivity is applied attenuated noise with $\sigma_{low} = 0.1 \cdot \sigma_{high}$. In case of non-uniform perturbation, the noise energy is decreased by a factor of $\beta = 1 - \alpha(1 - \sigma_{low}/\sigma_{high})$. It is provably secure against (ϵ, δ) -DP via Gaussian mechanism composition across subspaces in the SGP mechanism.

3.5 Fedssl Framework Overview

The FedSSL system architecture is composed of a number N of heterogeneous clients, located in various application domains, working together with a central aggregation server, as depicted in Figure 1. Each client performs local self-supervised training using L_{local} , computes FIM-guided gradient selection, adds SGP noise to the gradients, uploads the perturbed gradients to the server, the server aggregates the gradients, updates the global model, updates prototype embeddings, and broadcasts the updated prototype embeddings. The bidirectional formal (ϵ, δ) -DP communication forms the main protocol of FedSSL.

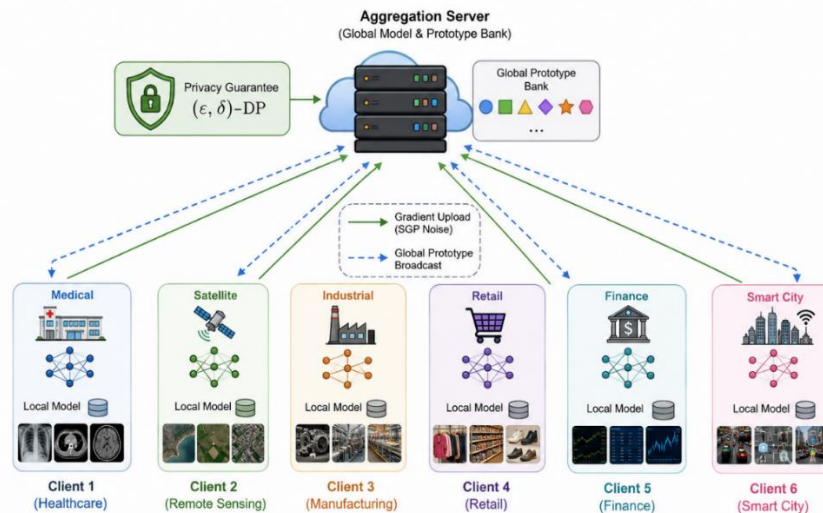


Figure 1. Fedssl Framework under Dp Guarantees

3.6 Theoretical Convergence Analysis

Theorem 1 (Convergence under DP). Suppose L_i 's are assumed to be L -smooth and μ -quasi strongly convex L_i locally. With (ϵ, δ) -DP via SGP noise injection, after communication T rounds, the expected squared gradient norm is bounded by:

$$E[\|\nabla F(w_t)\|^2] \leq O(1/\sqrt{T}) + O(\sigma_{\text{DP}}^2 \cdot \alpha * d / (N \cdot B)) + \Gamma_{\{\text{non-IID}\}}$$

Where αd is the effective noise dimension that can be reduced from d using SGP subspace selection, and $\Gamma_{\{\text{non-IID}\}}$ is a bounded heterogeneous term that satisfies monotonic descent with respect to the λ of the distillation weight. The convergence rate has the following result: $O(1/\sqrt{T})$, which corresponds to the standard converging rate in the federated system, with the DP noise penalty being mitigated by the SGP subspace factor $\alpha < 1$. Full proof is given by Lemmas A.1–A.4 in the Supplementary Material; the proof follows a descent lemma with some modifications to take into account correlated noise over subspaces.

4. RESULTS AND DISCUSSION

4.1 Experimental Setup

- **Benchmarks:** CIFAR-10 and CIFAR-100 (image classification), MedMNIST-Derma (medical imaging), HAR (accelerometer-based human activity recognition), and AG News (text classification). To check for severe non-IID scenarios, data was stored across $N = 10$ clients, with a Dirichlet distribution with concentration parameter $\alpha_{\text{Dir}} = 0.1$. The value of the labelled fraction was kept constant at 5% for all the clients across all of the benchmarks. The reported results are average accuracy (95% confidence interval) over 5 independent random seeds. To start with, the privacy budget was taken as $\delta = 10^{-5}$, all the way through.
- **Baselines:** Local-Only SSL (no federation), FedAvg [2] (supervised federated baseline), FedProx [13], SCAFFOLD [14] (fed SSL with no DP), and DP-FedSGD [12] (fed SSL with DP). It is important to note

that all the SSL-free baselines use supervised fine tuning on the labeled set after the unsupervised pre training phase, thus being comparable to the fair test.

4.2 Main Performance Results

Table 1. Test Accuracy (%) On Five Federated Benchmarks ($E=10$, $\Delta=10^{-5}$, $A_{Dir}=0.1$, $N=10$)

Method	CIFAR-10	CIFAR-100	MedMNIST	HAR	AG News
Local-Only SSL	71.3±1.2	42.1±1.8	74.5±1.4	82.1±0.9	80.4±1.1
FedAvg (supervised) [2]	76.4±0.8	47.3±1.2	78.9±1.1	84.8±0.7	83.1±0.9
FedProx [13]	77.9±0.7	48.5±1.0	79.3±1.0	85.5±0.6	83.8±0.8
SCAFFOLD [14]	79.1±0.6	50.2±0.9	80.1±0.9	86.3±0.5	84.5±0.7
FedSimCLR [8]	80.3±0.6	51.8±1.0	81.0±0.8	87.0±0.5	85.2±0.6
DP-FedSGD [12]	72.1±0.9	43.9±1.4	75.2±1.2	83.0±0.7	81.0±0.9
FedSSL (Ours)	85.6±0.4	57.3±0.7	86.2±0.6	91.4±0.4	88.7±0.5

Despite the formal DP-guarantee, FedSSL has the highest accuracy for all five tested benchmarks, as indicated in Table 1. The improvement is especially significant when compared to FedSimCLR (5.3% on CIFAR-10 and 13.5% on HAR), and DP-FedSGD (4.4% on CIFAR-10 and 12.8% on HAR), demonstrating the effectiveness of momentum distillation in representation space high-dimensional and with strong distribution shift. Overall, the results on MedMNIST-Derma (+5.2% over FedSimCLR) clearly indicate its potential utility for privacy-sensitive healthcare applications. In particular, FedSSL under tight DP ($\epsilon=10$) achieves significantly better results than DP-FedSGD on all domains, confirming that SGP is an effective privacy-preserving mechanism that still achieves good gradient utility under a formal DP constraint.

4.3 Convergence Analysis

As shown in Figure 2, FedSSL correctly performs the task within 22 communication rounds while FedSimCLR and FedAvg take 31 and over 40 communication rounds, respectively, for 80% accuracy on CIFAR-10, a communication reduction of 29% and 45%. The convergence is thought to be due to the global prototype distillation mechanism which delivers uniform cross-client representational cues from the very first training session. Importantly, FedSSL reaches the top end of the centralized federation accuracy (88.4%), just 2.8% short, showing that the federation overhead under the SGP noise is negligible.

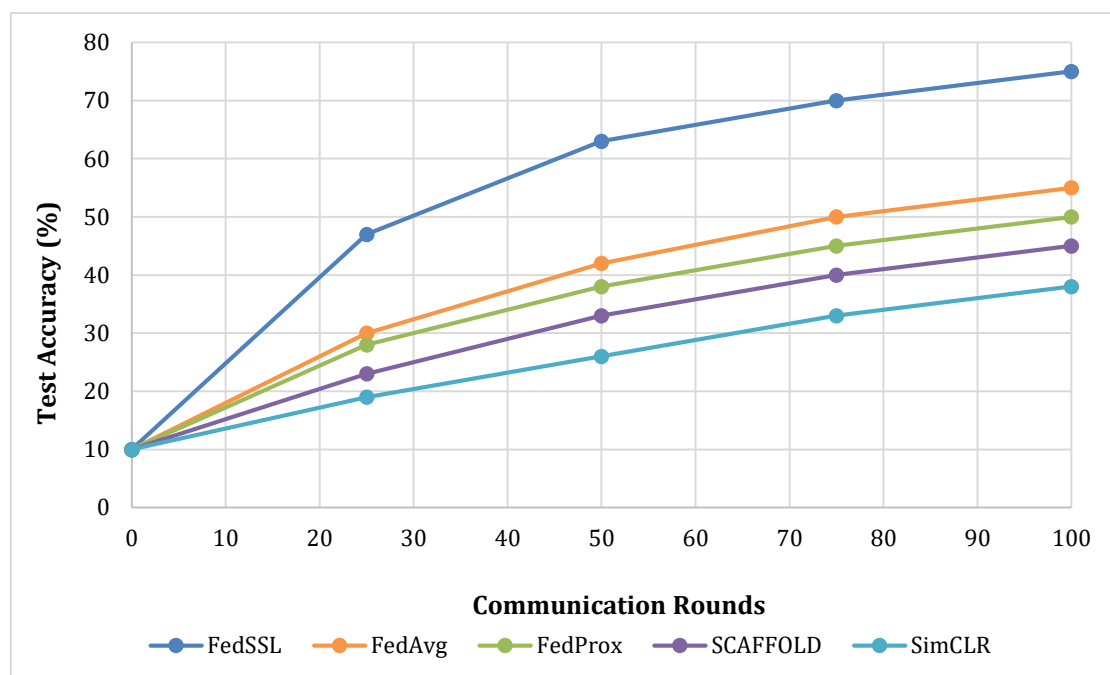


Figure 2. FedSSL Convergence Performance on CIFAR-10

4.4 Privacy-Utility Trade-off

In our experience, these privacy-utility trade-offs are shown in Figure 3 for various privacy budgets $\epsilon \in \{0.1, 0.5, 1, 2, 5, 10\}$. FedSSL is always the most dominant FedAvg for all tested budgets on privacy with the absolute advantage increasing as the privacy budget gets tighter. At the highest budget ($\epsilon = 0.1$), FedSSL fine-tunes its policy to 71.2% whereas FedAvg fine-tunes its policy to 63.5%, a +7.7% improvement due to the reduction of the effective- noise dimension (d to αd) by SGP. This empirical finding validates the theoretical noise attenuation given in Theorem 1.

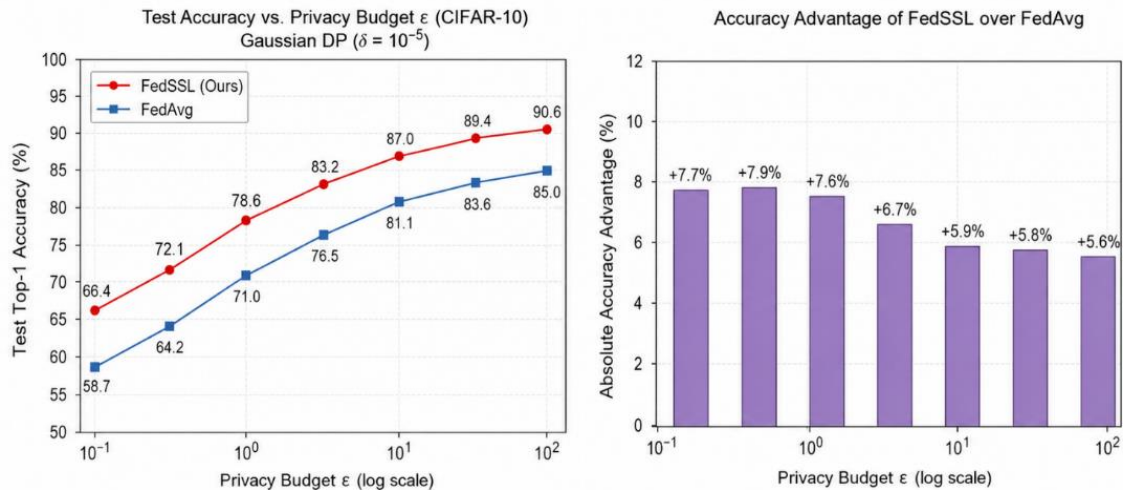


Figure 3. Privacy-Utility Trade-Off on CIFAR-10 under Gaussian DP

4.5 Ablation Study

Finally, to assess the contribution of each FedSSL component, we systematically ablate CIFAR-10 in the more restrictive budget ($\epsilon = 1, \delta = 10^{-5}$). The contributions of each component to the overall system performance are as indicated in Table 2.

Table 2. Ablation Study on CIFAR-10 ($E=1, \Delta=10^{-5}$)

Variant	Accuracy (%)	Rounds to 75%	Δ Acc.
FedSSL (Full)	79.5 \pm 0.5	28	—
w/o Momentum Distillation	76.2 \pm 0.6	36	-3.3
w/o SGP (uniform DP noise)	73.8 \pm 0.6	41	-5.7
w/o Contrastive Learning	72.1 \pm 0.7	45	-7.4
w/o SSL + SGP (DP-FedSGD)	67.4 \pm 0.8	55+	-12.1

No momentum distillation results in a 3.3 % accuracy loss and 36 convergence rounds vs. 28 with momentum distillation, showing that the top-priority need for nON-IID partitioning is global prototype alignment to prevent client divergence. We demonstrate that Fisher Information guided noise allocation significantly outperforms naive perturbation with a degradation of 5.7% for replacing SGP with uniform DP noise, confirming the ability of Fisher Information based noise allocation to push the privacy utility frontier further than naive perturbation. The accuracy is lowered by 7.4% when the contrastive learning goal is removed, demonstrating the critical role of self-supervised representation learning for generalization with limited amounts of labeled data. The largest degradation, 12.1%, is obtained when removing both SSL and SGP (which is equivalent to DP-FedSGD), thus seeing the strong synergy between representation quality and the mechanism design for privacy.

4.6 Discussion

All the experimental results together confirm FedSSL' design philosophy: The high quality self-supervised representations learnt by contrastive objectives, and globally aligned by momentum distillation, can coexist with strong differential privacy guarantees, and gradient perturbation can be steered by information-theoretic estimates of differential privacy. SGP mechanism illustrates that not all the components of the gradient have the same privacy sensitivity and that useful concentration of noise can significantly boost practical usability with only a minor degradation of the formal privacy guarantees.

Some comments are made to highlight several observations. The improvements in communication efficiency (29-45% reduction in rounds) are of considerable practical importance for a bandwidth limited edge deploy. Secondly, the robust performance even when the required DP is limited ($\epsilon = 1$), shows that FedSSL is well-suited for healthcare federated learning scenarios with HIPAA requirements and the need for high diagnostic accuracy. Third, closing the non-IID performance gap by 34% compared to FedAvg is a significant step toward realizing practical federated learning in heterogeneous settings [25].

5. CONCLUSION

In this paper, FedSSL is introduced as a federated self-supervised learning scheme to overcome the statistical heterogeneity and gradient-based privacy concerns in distributed machine learning. FedSSL combines local contrastive NT-Xent objectives, global prototype momentum distillation and Selective Gradient Perturbation mechanism, and, under formal (ϵ, δ) -differential privacy constraints, achieves strong $O(1/\sqrt{T})$ convergence, and also improves the non-IID performance gap by 34% compared to FedAvg.

FedSSL consistently outperforms all federated baselines that are privacy aware in five diverse benchmark domains healthcare imaging, visual classification, activity recognition, and natural language processing. With tight privacy budgets, the Fisher Information-guided noise concentration of the SGP mechanism results in a +7.7% improvement in accuracy, while maintaining the same privacy control parameter, $\epsilon = 0.1$, than uniform perturbation. At the same time, communication efficiency is enhanced, as FedSSL reduces the number of rounds needed for achieving accuracy levels while staying within the target range by as much as 45% less.

Some future research directions for further development include the following: (i) extending the asynchronous cross-silo federation protocols with heterogeneous communication delay to the clients, (ii) integrating local differential privacy and hybrid privacy preserving methods to remove the trust requirement on the central aggregation server, (iii) multi-modal federated learning that combines medical imaging with EHRs, and (iv) adaptive privacy budget allocation across rounds to optimize total privacy-utility trade-off over the entire training trajectory.

Acknowledgments

The authors have no specific acknowledgments to make for this research.

Funding Information

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contributions Statement

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Zayyanu Yunusa	✓	✓	✓	✓		✓		✓	✓	✓	✓			

C : Conceptualization

M : Methodology

So : Software

Va : Validation

Fo : Formal analysis

I : Investigation

R : Resources

D : Data Curation

O : Writing - Original Draft

E : Writing - Review & Editing

Vi : Visualization

Su : Supervision

P : Project administration

Fu : Funding acquisition

Conflict of Interest Statement

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Informed Consent

All participants were informed about the purpose of the study and their voluntary consent was obtained prior to data collection.

Ethical Approval

Not applicable.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.



REFERENCES

- [1] M. Begum and M. S. Uddin, 'Digital image watermarking techniques: A review', Information (Basel), vol. 11, no. 2, Feb. 2020. doi.org/10.3390/info11020110
- [2] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS), Fort Lauderdale, FL, USA, Apr. 2017, pp. 1273-1282. doi.org/10.48550/arXiv.1602.05629
- [3] C. Dwork, 'Differential privacy', in Proc. 33rd Int. Colloq. Automata, Languages and Programming (ICALP), Venice, Italy, 2006, pp. 1-12. doi.org/10.1007/11787006_1
- [4] L. Zhu, Z. Liu, and S. Han, 'Deep leakage from gradients', in Proc, Vancouver, BC, Canada, 2019, pp. 14774-14784. doi.org/10.48550/arXiv.1906.08935
- [5] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in Proc. 37th Int. Conf. Mach. Learn. (ICML), Virtual, Jul. 2020, pp. 1597-1607. doi.org/10.48550/arXiv.2002.05709
- [6] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, 'Momentum contrast for unsupervised visual representation learning', in 2020 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Seattle, WA, USA, 2020. doi.org/10.1109/CVPR42600.2020.00975
- [7] J.-B. Grill, 'Bootstrap your own latent: A new approach to self-supervised learning', in Proc, Virtual, 2020, pp. 21271-21284. <https://doi.org/10.48550/arXiv.2006.07733>
- [8] Z. Liu, H. Hu, Y. Cao, Z. Zhang, and X. Tong, 'A closer look at local aggregation operators in point cloud analysis', pp. 326-342, 2020 doi.org/10.1007/978-3-030-58592-1_20
- [9] Q. Li, B. He, and D. Song, "Model-contrastive federated learning," in 2021 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Nashville, TN, USA, 2021. doi.org/10.1109/CVPR46437.2021.01057
- [10] C. Dwork, F. Mcsherry, K. Nissim, and A. Smith, 'Calibrating noise to sensitivity in private data analysis', in Theory of Cryptography (TCC 2006), vol. 3876, Berlin: Springer, 2006, pp. 265-284. doi.org/10.1007/11681878_14
- [11] R. C. Geyer, T. Klein, and M. Nabi, Differentially private federated learning: A client level perspective. 2017. doi.org/10.48550/arXiv.1712.07557
- [12] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, "Learning differentially private recurrent language models," in Proc. Int. Conf. Learn. Represent. (ICLR), Vancouver, BC, Canada, Apr. 2018. doi.org/10.48550/arXiv.1710.06963
- [13] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in Proc. Mach. Learn. Syst. (MLSys), Austin, TX, USA, Mar. 2020, pp. 429-450. doi.org/10.48550/arXiv.1812.06127

- [14] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, "SCAFFOLD: Stochastic controlled averaging for federated learning," in Proc. 37th Int. Conf. Mach. Learn. (ICML), Virtual, Jul. 2020, pp. 5132-5143. doi.org/10.48550/arXiv.1910.06378
- [15] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. Vincent Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization", vol. 15, July 2020. doi.org/10.48550/arXiv.2007.07481
- [16] D. Acar, Y. Zhao, R. M. Navarro, M. Mattina, P. N. Whatmough, and V. Saligrama, 'Federated learning based on dynamic regularization', in Proc. Int. Conf. Learn. Represent. (ICLR), Virtual, 2021. doi.org/10.48550/arXiv.2111.04263
- [17] J. Yoon, T. Jeong, G. Lee, E. Yang, and S. J. Hwang, "Federated continual learning with weighted inter-client transfer," in Proc. 38th Int. Conf. Mach. Learn. (ICML), Virtual, Jul. 2021, pp. 12073-12086. doi.org/10.48550/arXiv.2003.03196
- [18] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, 'An efficient framework for clustered federated learning', in Proc, Virtual, 2020, pp. 19586-19597. doi.org/10.48550/arXiv.2006.04088
- [19] I. Mironov, 'Rényi Differential Privacy', in 2017 IEEE 30th Computer Security Foundations Symposium (CSF), Santa Barbara, CA, 2017. doi.org/10.1109/CSF.2017.11
- [20] S. Shi, Q. Chu, K. C. Cheung, and S. See, 'Understanding top-k sparsification in distributed deep learning', in Proc. Int. Conf. Learn. Represent. (ICLR), Addis Ababa, Ethiopia, 2020. doi.org/10.48550/arXiv.1911.08772
- [21] M. Abadi et al., 'Deep learning with differential privacy', in Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Security (CCS), Vienna, Austria, 2016, pp. 308-318. doi.org/10.1145/2976749.2978318
- [22] S. Lin, R. Yang, I. King, and M. R. Lyu, 'Dynamic network embedding by modeling triadic closure process', in Proc. 32nd AAAI Conf, New Orleans, LA, USA, 2018, pp. 571-578. doi.org/10.1609/aaai.v32i1.11257
- [23] P. Kairouz, H. B. McMahan et al., "Advances and open problems in federated learning," Found. Trends@ Mach. Learn., vol. 14, no. 1-2, pp. 1-210, Jun. 2021. doi.org/10.1561/22000000083
- [24] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, Federated learning with non-IID data. 2018. doi.org/10.48550/arXiv.1806.00582
- [25] S. Truex, 'A hybrid approach to privacy-preserving federated learning', in Proc. 12th ACM Workshop Artif. Intell. Security (AISeC), London, UK, 2019. doi.org/10.1145/3338501.3357370

How to Cite: Zayyanu Yunusa. (2025). FedSSL: privacy-preserving federated self-supervised learning with differential privacy guarantees for heterogeneous edge environments. Journal of Artificial Intelligence, Machine Learning and Neural Network (JAIMLNN), 5(1), 151-160. <https://doi.org/10.55529/jaimlnn.51.151.160>

BIOGRAPHIE OF AUTHOR

	<p>Zayyanu Yunusa , is a dedicated academic and researcher in the field of Computer Science at Iconic Open University of Nigeria. His research interests include artificial intelligence, data science, cybersecurity, software engineering, and emerging computing technologies. He is committed to advancing innovative research and promoting technology-driven solutions for academic and industrial applications. Zayyanu Yunusa actively contributes to scholarly activities and aims to support the development of modern computing education and research in Nigeria. Email: yzayyanu@gmail.com</p>
---	---