AFRC: Adaptive Responsible Compression for Federated Learning under Data Heterogeneity

Extended Abstract

Adarsh Valoor
University of Southampton
Southampton, United Kingdom
adarsh.valoor@soton.ac.uk

Sarvapali D. Ramchurn
University of Southampton
Southampton, United Kingdom
sdr1@soton.ac.uk

ABSTRACT

We present AFRC (Adaptive Federated Responsible Compression), a coordination mechanism for multiagent federated learning that jointly regulates model compression, cross-agent equity, and differential privacy under statistical heterogeneity. AFRC introduces two feedback controllers: a proportional-integral (PI) fairness controller that dynamically adjusts per-round fairness pressure to drive equitable agent outcomes, and a budget-aware privacy controller that schedules the DP noise multiplier to approximately meet a global (ε, δ) target while preserving late-stage utility. Across K=100 agents on CIFAR-10 and Shakespeare with severe non-IID partitions (α=0.1), AFRC achieves 3–7% higher average accuracy than the strongest non-private baseline, while reducing inter-agent accuracy variance by over 30% compared to FedAvg, at 90% sparsity and a fixed ε=5 budget. We provide convergence guarantees for a simplified variant and show that adaptive mechanism-based coordination is essential to balance utility, equity, privacy, and efficiency in decentralised multiagent learning.

CCS CONCEPTS

• **Computing methodologies → Multi-agent systems**.

KEYWORDS

Adaptive Control, Coordination Mechanisms, Differential Privacy, Fairness, Federated Multiagent Learning

ACM Reference Format:
Adarsh Valoor and Sarvapali D. Ramchurn. 2026. AFRC: Adaptive Responsible Compression for Federated Learning under Data Heterogeneity: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/OIBA2819

1 INTRODUCTION

Federated Learning (FL) [6] enables decentralised multiagent systems to collaboratively train shared models without pooling private data. In practice, deployments must satisfy four coupled objectives: **accuracy**, **communication efficiency**, **cross-agent fairness**, and **differential privacy (DP)**. These objectives are especially difficult to balance under non-IID data, where compression can disproportionately harm minority agents [1], and fixed DP noise degrades both accuracy and fairness [2].

The dominant practice is a *static* compromise: fixed compression ratio, fairness coefficient, and noise level. This is brittle — configurations that stabilise the turbulent early phase of training are rarely optimal for fine-tuning. We propose AFRC, which replaces static parameters with two online feedback controllers that self-regulate the balance between all four objectives across training rounds.

Contributions. (1) A coordination mechanism for responsible FL that integrates dual feedback controllers into the standard aggregation loop. (2) A complete algorithmic specification including PI fairness control with anti-windup, RDP-based budget-aware DP scheduling via Poisson subsampling, and structured pruning. (3) A convergence guarantee for a simplified variant, decomposing the optimisation neighbourhood into a fairness-bias term and a DP/variance term. (4) Empirical evaluation on CIFAR-10 and Shakespeare with 100 agents, severe heterogeneity, and full per-agent outcome reporting.

2 FRAMEWORK AND ALGORITHM

Setup. Consider K agents with private datasets $\{D_k\}$ where $D_k \sim P_k$, $P_k \neq P_j$. Each round t, the server samples $S_t \subset [K]$ with participation rate $q = |S_t|/K$ and aggregates using data-proportional weights:

$$w_{t+1} = w_t + \sum_{k \in S_t} \frac{|D_k|}{\sum_{j \in S_t} |D_j|} \tilde{\Delta}_t^k, \tag{1}$$

where $\tilde{\Delta}_t^k$ is the clipped and privatised local update (defined below). Each agent minimises a fairness-regularised local objective:

$$\min_{w^k} \left[F_k(w^k) + \lambda_{F,t} \|w^k - w_t\|_2^2 + \mathcal{R}(w^k) \right], \tag{2}$$

with global adaptive fairness weight $\lambda_{F,t} \geq 0$ and compressibility regulariser $\mathcal{R}$.

PI Fairness Controller. Let $\bar{a}_t = \frac{1}{|S_t|} \sum_{k \in S_t} a_t^k$ denote the mean per-agent validation accuracy in round t. A rolling dispersion estimator tracks cross-agent outcome variance:

$$v_t = (1 - \gamma)\, v_{t-1} + \gamma \cdot \frac{1}{|S_t|} \sum_{k \in S_t} \left(a_t^k - \bar{a}_t\right)^2, \tag{3}$$

with exponential smoothing parameter $\gamma \in (0, 1)$. Defining error $e_t = v_t - v^*$ relative to a target dispersion band $v^* \geq 0$, the PI update

---

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/OIBA2819

with conditional-integration anti-windup is:

$$s_t \leftarrow s_{t-1} + e_t \cdot \mathbf{1}\left\{\lambda_{F,t} \in (0, \lambda_{\max})\right\}, \tag{4}$$

$$\lambda_{F,t+1} \leftarrow \Pi_{[0,\lambda_{\max}]}\left(\lambda_{F,t} + k_p e_t + k_i s_t\right). \tag{5}$$

The indicator freezes integration at the actuator limits ($\lambda_{F,t} = 0$ or $\lambda_{F,t} = \lambda_{\max}$), preventing integral windup. When dispersion rises, $\lambda_{F,t}$ increases, tightening the drift penalty and pulling minority-agent updates toward the global model; as v_t enters the target band, $\lambda_{F,t}$ relaxes.

Budget-aware DP Scheduler. Each agent clips its update to ℓ_2 norm C and adds calibrated Gaussian noise:

$$\tilde{\Delta}_t^k = \mathrm{clip}(\Delta_t^k, C) + \mathcal{N}\left(0, \sigma_t^2 C^2 I\right). \tag{6}$$

Given remaining privacy budget $\varepsilon_{\mathrm{rem}}$ at the start of round t, the scheduler selects the smallest feasible noise multiplier σ_{t+1} by solving the constrained problem:

$$\begin{aligned} &\min_{\sigma \ge 0} \ \sigma \\ &\text{s.t.} \quad \min_{\alpha \in \mathcal{A}} \left[\varepsilon_\alpha(\sigma, q) + \frac{\log(1/\delta)}{\alpha - 1}\right] \le \frac{\varepsilon_{\mathrm{rem}}}{T - t}, \end{aligned} \tag{7}$$

where $\varepsilon_\alpha(\sigma, q)$ is the RDP guarantee of the subsampled Gaussian mechanism at order α under Poisson subsampling [8], and $\mathcal{A}$ is a finite grid of RDP orders. The inner min converts the per-round RDP guarantee to a $(\cdot, \delta)$-DP cost via the standard RDP-to-(ε, δ) conversion [8]; the outer minimisation over σ is solved by Brent's method. Allocating $\varepsilon_{\mathrm{rem}}/(T-t)$ uniformly across remaining rounds front-loads noise (higher σ early, lower σ late), preserving utility in the fine-tuning phase.

Structured Pruning. After each aggregation, the server applies magnitude-based structured pruning following a smooth schedule $s_t \nearrow s_{\max}$ (default 90%), reducing uplink/downlink communication without destabilising training.

3 CONVERGENCE ANALYSIS

Under standard assumptions — ℓ-smooth global loss, bounded stochastic gradient variance $\mathbb{E}\|\nabla F_k\|^2 \le S^2$, bounded data heterogeneity ζ^2, and bounded controller outputs ($\lambda_{F,t} \le \lambda_{\max}$, $\sigma_t \le \sigma_{\max}$) — we prove:

Theorem 3.1 (Convergence to a neighbourhood). *For the simplified AFRC update (clipping, local DP, data-proportional aggregation, fairness regularisation; pruning excluded) with learning rate $\eta \le 1/\ell$, letting $B_{\max}^2 = \max_t \mathbb{E}\|B_t\|^2$ bound the aggregation bias:*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla F_{\mathrm{global}}(w_t)\right\|^2 \ \le \ \frac{2(F_0 - F^*)}{T\eta} + \underbrace{B_{\max}^2}_{\textit{fairness bias}} + \underbrace{\ell\,\eta\,G^2}_{\textit{DP + stoch. var.}}, \tag{8}$$

where $G^2 \le S^2 + d\,\sigma_{\max}^2 C^2$, $F_0 = F_{\mathrm{global}}(w_0)$, and $F^* = \inf_w F_{\mathrm{global}}(w)$. Setting $\eta = \sqrt{(F_0 - F^*)/(\ell G^2 T)}$ gives an $O(1/\sqrt{T})$ convergence rate to the neighbourhood.

The PI fairness controller shrinks $B_{\max}^2$ by contracting extreme client drift; the DP scheduler bounds the noise term by capping $\sigma_{\max}$ via the remaining budget. Anti-windup and projection in Eqs. (4)–(5) justify the bounded-adaptation assumption.

4 EXPERIMENTS

Setup. K=100 agents, participation rate q=0.1, T=200 rounds, Dirichlet non-IID with $\alpha \in \{0.5, 0.1\}$ (severe), (ε=5, δ=10^{-5}), $s_{\max}$=90%. Datasets: CIFAR-10 (ResNet-18) and Shakespeare (char-LSTM). All results are mean ± sd over 5 seeds.

Baselines. FedAvg [6], FedAvg+Pruning, q-FFL [5], DP-FedAvg [7], FedProx [4], SCAFFOLD [3], Top-k sparsified [10], FedPAQ [9]. Ablations remove each controller independently.

Results. Table 1 summarises end-of-training outcomes under severe heterogeneity (α=0.1). AFRC achieves the highest average accuracy on both datasets (78.0% on CIFAR-10, 59.3% on Shakespeare) while simultaneously attaining the lowest variance and highest 10th-percentile accuracy, at a matched privacy budget ε=5.

Table 1: Performance after 200 rounds (α=0.1, ε=5), selected baselines. Best result bold.

Method	Avg. Acc.%	Var.	P10%	ε
CIFAR-10 / ResNet-18				
FedAvg	71.2	0.089	52.0	–
q-FFL	70.5	0.055	58.4	–
DP-FedAvg	65.3	0.102	45.0	5.0
SCAFFOLD	73.0	0.066	57.5	–
AFRC	**78.0**	**0.052**	**65.2**	**5.0**
Shakespeare / char-LSTM				
FedAvg	55.1	0.075	38.2	–
q-FFL	54.8	0.041	41.5	–
DP-FedAvg	48.9	0.091	30.3	5.0
SCAFFOLD	56.0	0.057	41.0	–
AFRC	**59.3**	**0.035**	**48.8**	**5.0**

Ablations confirm each controller's necessity: removing the PI controller raises variance by +44% on CIFAR-10; removing the DP scheduler drops average accuracy by ~6% due to excess late-stage noise. The combined gain exceeds the sum of individual contributions, evidencing synergy — fairness control improves the effective signal-to-noise ratio under DP, and the reduced late-stage noise in turn narrows per-agent accuracy disparity. Communication cost matches FedAvg+Pruning (identical pruning schedule) while reaching target accuracy with fewer total MB than update-sparsification baselines.

5 CONCLUSION

AFRC demonstrates that the four-way tension between utility, equity, privacy, and efficiency in heterogeneous multiagent FL cannot be resolved by static hyperparameter choices. Dual feedback controllers — a PI fairness regulator and a budget-aware DP scheduler — jointly adapt online to training dynamics, yielding consistent improvements across tasks, heterogeneity regimes, and ablation conditions. Future work will extend the convergence guarantee to time-varying controllers and non-smooth pruning, incorporate per-agent DP budgets, and evaluate on transformer-based models with adapter training.

ACKNOWLEDGMENTS

This research was suported by the Engineering and Physical Sciences Research Council (grant number EP/Y009800/1), through funding from Responsible Ai UK.

REFERENCES

[1] Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. 2020. Characterising Bias in Compressed Models. *arXiv preprint arXiv:2010.03058* (2020). https://arxiv.org/abs/2010.03058

[2] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, and et al. 2021. Advances and Open Problems in Federated Learning. *Foundations and Trends in Machine Learning* 14, 1-2 (2021), 1–210. https://doi.org/10.1561/2200000083

[3] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. 2020. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 5132–5143. https://proceedings.mlr.press/v119/karimireddy20a.html

[4] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *Proceedings of Machine Learning and Systems (MLSys)*, Vol. 2. 429–450. https://proceedings.mlsys.org/paper/2020/hash/1f5fe83998a09396ebe6477d9475ba0c-Abstract.html

[5] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. 2020. Fair Resource Allocation in Federated Learning. In *International Conference on Learning Representations (ICLR)*. https://openreview.net/forum?id=ByexElSYDr

[6] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (Proceedings of Machine Learning Research, Vol. 54)*. PMLR, 1273–1282. https://proceedings.mlr.press/v54/mcmahan17a.html

[7] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *International Conference on Learning Representations (ICLR)*. https://openreview.net/forum?id=BJ0hF1Z0b

[8] Ilya Mironov. 2017. Rényi Differential Privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, 263–275. https://doi.org/10.1109/CSF.2017.11

[9] Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. 2020. FedPAQ: A Communication-Efficient Federated Learning Method with Periodic Averaging and Quantization. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS) (Proceedings of Machine Learning Research, Vol. 108)*. PMLR, 2021–2031. https://proceedings.mlr.press/v108/reisizadeh20a.html

[10] Sebastian U. Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. 2018. Sparsified SGD with Memory. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*. 4452–4463. https://papers.nips.cc/paper/2018/hash/b440509a0106086a67bc2ea9df0a1dab-Abstract.html