

Omni-Governance Framework (OGF): A Tiered Approach to GDPR-Compliant, Bias-Resilient AI in Decentralized Ecosystems

1st Suha Afaneh
Department of Cybersecurity
Zarqa University
Zarqa, Jordan

2nd Rami Almatarneh
Department of Cybersecurity
Zarqa University
Zarqa, Jordan

Abstract

The rapid evolution of Web 4.0, which is fundamentally based on a decentralized, AI-powered internet, demands governance frameworks that harmonize innovation with ethical imperatives. Current solutions for privacy preservation and algorithmic fairness in decentralized ecosystems remain fragmented, struggling to balance regulatory compliance (GDPR, EU AI Act) with technical scalability and bias resilience. In this paper we present the Omni-Governance Framework (OGF), a tiered architecture designed to address three key challenges: (1) the tradeoff between privacy and fairness in federated learning, (2) regulatory inconsistencies across jurisdictions, and (3) the growing vulnerability of classical encryption in a post-quantum world.

The OGF architecture integrates three interconnected tiers: the first tier (T1) enforces GDPR compliance through Rényi Differential Privacy (RDP) and a scalable blockchain-based consent mechanism that mathematically links privacy budgets (ϵ) to fairness degradation (ΔF) to measure how much privacy is preserved (ϵ) with how much fairness might be affected (ΔF), helping to strike a thoughtful balance between the two; the second tier (T2) mitigates bias through fairness-constrained synthetic data generation and edge AI deployment, reducing latency by optimizing computational costs $\tau = C(M)/P_{\text{edge}}$; and finally, the third tier (T3) introduces quantum-safe governance with lattice-based homomorphic encryption $\text{Enc}(VW) = As + e \bmod q$ and Decentralized Autonomous Ethics Committees (DAECs) that use quadratic voting $\text{Vote Cost} = \sum w_i^2$ to democratize decision-making. OGF's cross-tier synergies resolve critical gaps in existing frameworks like FATE and PySyft, for example, synthetic data from Tier 2 enhances federated dataset diversity in Tier 1, while neuromorphic auditors in Tier 3 detect bias drift in real time via spiking neural networks $B = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N s_i(t)$. The modular structure of the proposed framework enables incremental adoption, depending on an organization's readiness and goals, in line with regulatory maturity, in addition, its quantum-resistant design prevents emerging threats to decentralized AI. While challenges remain, such as a 15% increase in latency at Tier 3 due to encryption overhead and ethical concerns associated with tokenized redress, OGF offers a unified framework to bridge the gap between current regulatory requirements with future technological limitations. This work presents a scalable blueprint for ethical AI governance in the Web 4.0 era, ensuring privacy preservation, bias mitigation, and democratized accountability.

Keywords: *Decentralized AI Governance, GDPR Compliance, Algorithmic Bias Mitigation, Quantum-Resistant Encryption, Ethical Machine Learning.*

1. Introduction

The emergence of Web 4.0 promises to completely revolutionize industries through the seamless integration of artificial intelligence (AI), blockchain, and edge computing, due to its decentralized, intelligent, and highly connected version of the Internet. Characterized by autonomous systems, decentralized data ecosystems, and human-machine collaboration, Web 4.0 envisions a world where smart cities, personalized healthcare, and democratized finance operate within interconnected yet privacy-preserving frameworks [1]. Even so, the Web 4.0 vision still has particular challenges in balancing innovation and ethical imperatives. As AI-driven systems become ubiquitous in Web 4.0, two critical issues loom large: compliance with stringent regulations like the EU's General Data Protection Regulation (GDPR) and the mitigation of algorithmic bias in decentralized data ecosystems [2]. These challenges are exacerbated by the inherent tension between AI's reliance on vast, distributed datasets and regulatory mandates for data minimization, transparency, and user control [3].

Current governance approaches in decentralized AI systems are still fragmented. Frameworks such as Federated learning, for example, give priority to privacy by design, enabling model training on distributed nodes without centralizing sensitive data, but often neglect to address systemic bias propagation [4]. Blockchain-based consent management tools effectively enhance transparency, but they often face difficulties in reconciling immutability with the “right to be forgotten” enshrined in the General Data Protection Regulation (GDPR’s) [5]. While some fairness aware AI toolkits such as IBM’s AIF360 aim to mitigate bias in algorithmic and human decision making, these frameworks are not developed with decentralized infrastructures in mind [6]. This disjunction creates a fragmented environment where organizations end up piecing together various solutions, usually focusing on compliance after the fact rather than building ethical principles into their systems from the start, which will lead to hard to ignore consequences, for example, facial recognition technologies used in smart cities have shown error rates exceeding 20% for marginalized communities [7], and decentralized finance (DeFi) platforms inadvertently encode socioeconomic prejudices via poorly designed credit-scoring algorithms [8].

The EU AI Act of 2024 compounded these challenges by classifying high-risk AI systems, such as those used in healthcare or recruitment, by imposing strict bias audits along with human oversight [9]. However, the Act's enforcement mechanisms are still in their infancy, especially in decentralized contexts where data ownership is fragmented across devices, institutions, and jurisdictions [10]. The global scalability of Web 4.0 further complicates compliance as cross-border data flows clash with regional regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) [11]. In addition, the looming threat of quantum computing could compound these issues by making classical cryptographic methods obsolete, threatening the privacy foundations of current decentralized systems [12].

In this paper, we propose the Omni-Governance Framework (OGF), a tiered, adaptive architecture designed to harmonize privacy, fairness, and innovation in Web 4.0 ecosystems. Unlike prior frameworks, OGF adopts a modular approach, that enables organizations to incrementally adopt governance capabilities aligned with their technical maturity and ethical aspirations.

On its base layer (Tier 1), OGF combines GDPR-compliant federated learning with blockchain-enabled consent management, providing data minimization and auditability for distributed AI workflows [13]. Tier 2 additionally overlays sophisticated synthetic data oracles and edge AI deployment to prevent bias propagation by real-time preprocessing and automated compliance checks [14]. Finally, Tier 3 is leading the way in solutions like quantum-resistant homomorphic encryption and Decentralized Autonomous Ethics Committees (DAECs) with quadratic voting to decentralize world-scale AI system governance [15].

The innovation of OGF’s lies in its cross-tier synergies. For instance, synthetic data generated in Tier 2 enhances the diversity of federated learning datasets in Tier 1, and reduce bias at its source [16].

Meanwhile, Level 3 neuromorphic validators (energy-efficient, optimized neural networks deployed on edge devices) detect bias drift in real time, triggering DAEC-led governance adjustments [17]. This closed-loop system not only ensures compliance with evolving regulations like the EU AI Act, but also ensures that Web 4.0 ecosystems are protected from quantum-era threats [18].

The rapid expansion of Web 4.0 technology emphasizes the need for such a framework. By 2030, more than 75% of companies are expected to operate in decentralized environments, yet less than 15% of them have implemented comprehensive governance strategies that combine privacy and fairness [19, 20].

Yet, OGF is not without some challenges. For example, the computational overhead of quantum-safe encryption may threaten to slow down Tier 3 workflows, while DAEC governance models must guard against populist biases that could overshadowing expert consensus [21, 22]. In addition, ethical concerns also arise around tokenized compensation mechanisms, which, while innovative, risk commodifying redress for algorithmic harm [23]. Nevertheless, OGF represents a critical step toward a Web 4.0 future where innovation and ethics coexist. By uniting foundational compliance, advanced technical safeguards, and speculative governance models, it provides a blueprint for trustworthy AI in an increasingly decentralized world, one where privacy is preserved, biases are preempted, and accountability is democratized.

2. Literature Review

The regulation of decentralized AI systems in Web 4.0 ecosystems has become an important topic of research, propelled by the need for regulatory compliance and ethical accountability. Initial researches on decentralized AI, mainly focused on technical innovations, such as federated learning architectures that enable collaborative model training without centralized data aggregation [1].

Although these developments dealt with privacy to a certain degree, they often overlooked the systemic biases embedded in distributed datasets. For example, foundational work in federated learning demonstrated low privacy

risks but recognized that local data biases may propagate into global models, which exacerbates disparities in applications such as healthcare diagnostics [4]. This limitation emphasized the need for frameworks that would combine privacy protection and bias prevention, a gap subsequently explored by fairness toolkits designed for centralized AI, which lack adaptation to decentralized environments [6].

The intersection of GDPR compliance and decentralized AI gained prominence as Web 4.0 infrastructures expanded. Blockchain-based personal data management systems were suggested for providing higher transparency, although early designs lacked the complete reconciling of blockchain immutability with the "right to be forgotten" of GDPR [3]. Subsequent studies attempted to resolve this conflict by designing mutable blockchain layers for consent revocation, though scalability remained a challenge in large-scale networks [5]. Simultaneously, regulatory innovation in the shape of the EU AI Act placed strict obligations on high-risk AI systems, such as bias auditing and human oversight, a turn of events that brought to light that the available toolkit was insufficient for decentralized environments [9]. For instance, the failure of fairness-aware algorithms that operate on static datasets to consider dynamic feedback loops in transactional Web 4.0 systems, such as decentralized finance (DeFi) platforms [13].

Algorithmic bias in decentralized systems is also well-documented, especially in applications with social effects on marginalized groups. Pioneering studies on racial and gender gaps in facial recognition technology has shown that when training data lacks diversity, this in general leads to biased results, and it becomes even more pronounced in decentralized systems, where isolated data sources tend to mirror the specific biases of their region or institution [7]. Earlier research broke down the sources of bias into three main types: data, design, and interaction. It also pointed out that decentralized systems tend to make interaction bias worse, mainly because feedback loops in these systems can reinforce and amplify existing issues over time [8]. Recent research investigated synthetic data generation as a solution and showed that fairness-constrained generative adversarial networks (GANs) can be utilized to improve dataset diversity in federated learning [14]. Yet findings also alerted to synthetic data's potential to duplicate hidden biases if validation processes are flawed—a vulnerability intensified in Web 4.0's heterogeneous data environments [16].

Efforts to harmonize General Data Protection Regulation (GDPR) compliance with AI fairness are increasingly turning to edge computing and neuromorphic technologies. Studies demonstrated that spiking neural networks (SNNs) deployed on edge devices could detect bias in real time with high accuracy while consuming minimal power, a breakthrough for energy-efficient governance in IoT-driven Web 4.0 networks [17]. In the meantime, quantum-resistant homomorphic encryption was designed to protect federated learning against attacks from quantum computers, but without considering bias propagation in encrypted gradients [12]. These technical advancements, as promising as they were, also worked in silos, ignoring Web 4.0's global span and integrated governance that it demands.

Decentralized AI governance models have evolved from centralized audits to decentralized autonomous organizations (DAOs). Quadratic voting proposals in these organizations have sought to democratize decision-making, allowing stakeholders to proportionally influence fairness thresholds and privacy budgets [15]. This was, nevertheless, criticized as susceptible to manipulation by influential actors, a reflection of more general anxieties regarding power asymmetries in decentralized environments [22]. Regulatory research called for adaptive frameworks that directly relate to developing law such as the EU AI Act, even if models were not technically specific to Web 4.0 decentralized infrastructures [10].

The ethical implications of AI redress mechanisms have also garnered attention. Tokenized compensation systems on blockchain ledgers have been proposed to automate compensation for biased outcomes, but this approach risks reducing ethical accountability to transactional exchanges [18]. On the other hand, criticism warned against commodifying algorithmic harm and have instead called for structural reforms in AI governance, a perspective that aligns with calls for AGI-ready ethical frameworks that proactively enshrine fairness and privacy [23].

In spite of these evolutions, serious gaps remain. First, current solutions tend to tackle GDPR compliance or bias mitigation separately, without consideration of their mutual dependency in decentralized ecosystems [2]. Second, synthetic data solutions lack robust validation protocols for Web 4.0's dynamic environments, threatening to amplify bias [21]. Third, quantum-resistant technologies remain far from fairness objectives, leaving future systems vulnerable to both privacy breaches and ethical failures [11], and finally, governance models like DAOs prioritize decentralization over expert oversight, which can undermine accountability in high-risk areas such as healthcare [20].

Against this backdrop, the Omnibus Governance Framework (OGF) emerged, bringing together disparate innovations into a multi-tiered architecture, so OGF can bridge the divide between current capabilities and future challenges by integrating foundational privacy techniques, advanced bias mitigation, and visionary governance. OGF cross-tier synergies, such as using synthetic data to enhance fairness will offers a more unified alternative to the often-fragmented solutions seen in the past [13]. Moreover, OGF's quantum-safe design proactively secures Web 4.0 ecosystems against emerging threats while maintaining backward compatibility with classical infrastructures [12].

3. Omni-Governance Framework (OGF): A Unified Architecture for Ethical AI in Web 4.0

The Omni-Governance Framework (OGF) addresses critical challenges in Web 4.0 ecosystems, which mainly depends on decentralized AI systems that demand rigorous privacy preservation, algorithmic fairness, and adaptive governance. Existing frameworks like FATE (federated learning) and PySyft (privacy-preserving ML) focus narrowly on privacy or federated workflows but fail to integrate fairness mechanisms, regulatory compliance, and future-proof governance [24]. OGF bridges these gaps through a tiered architecture that harmonizes three evolutionary layers, enabling organizations to scale from foundational compliance to visionary AGI-ready systems.

3.1 Core Challenges and Rationale

Web 4.0's decentralized AI systems face four still unsolved challenges:

1. **Privacy-Fairness Tradeoffs:** Federated learning (FL) minimizes data centralization but amplifies bias propagation (e.g., skewed local datasets in healthcare models [25]). OGF solves this via Rényi differential privacy (RDP), which measures fairness degradation ΔF under a privacy budget ϵ :

$$\Delta F = \frac{\partial F}{\partial \epsilon} \cdot \Delta \epsilon$$

where F is fairness.

2. **Regulatory Fragmentation:** Static tools like IBM AIF360 lack the ability to adapt to evolving regulations, such as the European Union's AI law (EU AI Act). OGF's regulatory ontology engine dynamically matches Article 22 of the General Data Protection Regulation (GDPR) to AI Act risk levels using \mathcal{R} semantic rules:

$$\mathcal{R} \vdash \text{GDPR}_{\text{Art.22}} \rightarrow \text{AIAct}_{\text{Risk-Level}}$$

3. **Quantum Vulnerabilities:** Classical encryption methods, such as AES-256 are expected to be vulnerable to quantum attacks, and to address this issue, OGF takes a more future-proof approach by using lattice-based homomorphic encryption to protect gradients in federated learning systems [25].
4. **Centralized Governance:** Current models (e.g., centralized audits) conflict with Web 4.0's decentralized ethos. OGF's **Decentralized Autonomous Ethics Committees (DAECs)** use quadratic voting to democratize governance:

$$\text{Vote Cost} = \sum_{i=1}^n w_i^2$$

Where w_i voter weight.

3.2 OGF Components and Architecture

The Omni-Governance Framework (OGF) is a multi-tiered architecture which is designed to align privacy, fairness, and governance in decentralized AI systems that support Web 4.0. It consists of three layers (see table 1): the data layer (T1), which promotes compliance with the General Data Protection Regulation (GDPR) standards through federated learning with Reni's differential privacy (RDP) and blockchain-based consent management; the processing layer (T2), which mitigates bias through edge AI deployment and fairness-constrained synthetic data generation; and the governance layer (T3), which promotes future systems through quantum-secure cryptography, neuro-bias auditors, and decentralized autonomous ethics committees (DAECs) [26]. Redress mechanisms, a self-healing model updates and blockchain-based compensatory tokens, close the accountability loop, ensuring biased outcomes trigger automated corrections and user compensation [27].

Table 1: OGF Tier Alignment with GDPR and the EU AI Act

OGF Tier	GDPR Compliance	EU AI Act Alignment
Tier 1	Art. 5 (Data minimization)	Annex III (High-risk systems) [9]
Tier 2	Art. 22 (Automated decisions)	Art. 13 (Transparency) [9]
Tier 3	Art. 25 (Privacy by design)	Art. 52 (Post-market monitoring) [9]

The OGF architecture integrates three tiers (see figure 1), each of which designed to address distinct technical and ethical challenges, while also working together to create meaningful synergy across the layers:

1. Tier 1: Foundational Layer (Privacy Compliance)

- **Federated Learning with Rényi DP:**

For N nodes, local model updates ΔW_i are aggregated with RDP noise η :

$$W_{global} = \frac{1}{N} \sum_{i=1}^N (\Delta W_i + \eta), \quad \eta \sim \mathcal{N}(0, \sigma^2 I)$$

where σ^2 ensures (ϵ, δ) -DP guarantees [28].

- **Blockchain Consent Management:**

- Consent records c_i are hashed as $H(c_i)$ and stored on a **Hyperledger Fabric** blockchain, with smart contracts automating GDPR's "right to erasure" via zero-knowledge proofs [29].

2. Tier 2: Advanced Layer (Bias Mitigation)

- **Edge AI Deployment:**

- Models deployed on edge devices \mathcal{E} optimize latency τ and energy \mathcal{P} :

$$\tau = \frac{\mathcal{C}(M)}{\mathcal{P}_{edge}} \quad (\mathcal{C}(M))$$

where $\mathcal{C}(M)$ computational cost of model M .

- **Synthetic Data Oracles:**

- Fairness-constrained GANs generate synthetic data \mathcal{S} with a fairness loss \mathcal{L}_f :

$$\min_G \max_D \mathcal{L}_{GAN} + \lambda \mathcal{L}_f, \quad \mathcal{L}_f = \sum_{g \in \mathcal{G}} |\mathbb{E}[\hat{y}_g] - \mathbb{E}[\hat{y}]|$$

where \mathcal{G} denotes protected groups [30].

3. Tier 3: Visionary Layer (Future-Proof Governance)

- **Quantum-Safe Homomorphic Encryption (QS-HE):**

- FL gradients ∇W are encrypted using **CRYSTALS-Kyber** lattice cryptography:

$$\text{Enc}(\nabla W) = As + e \text{ mod } q$$

where A public matrix and s secret key.

- **Neuromorphic Auditors:**

- Spiking neural networks (SNNs) on Intel Loihi chips detect bias \mathcal{B} via spike-counting:

$$\mathcal{B} = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N s_i(t)$$

where $s_i(t)$ binary spike output.

- **DAEC Governance:**

- Stakeholders vote on fairness thresholds δ using quadratic voting, with weights w_i proportional to token holdings t_i :

$$w_i = \sqrt{t_i} \quad (\text{ensuring quadratic cost proportionality [15]})$$

DAEC governance uses quadratic voting to democratize fairness thresholds. To mitigate risks (e.g., token concentration), OGF implements:

- **Sybil resistance:** Proof-of-stake validation limits duplicate accounts [15].
- **Expert oversight:** 20% voting weight reserved for ethics auditors [22].

The unique design of OGF integrates privacy-by-design, dynamic fairness, and democratic governance into a cohesive workflow, making it a stark contrast to siloed frameworks like FATE or PySyft, and ideally suited to Web 4.0's decentralized and interconnected ecosystems. Moreover, its modularity allows for easy incremental adoption, which is an important advantage for organizations considering transitioning from legacy systems to Web 4.0 [31]. Quantum-safe encryption and neuromorphic auditors prevent emerging threats, while DAEC governance aligns with decentralized decision-making norms, ensuring compliance with evolving regulations like the EU AI Act [15]. One of the framework's key strengths is its closed-loop accountability system, when bias is detected, it automatically triggers corrective action. This approach helps overcome one of federated learning's biggest challenges: balancing privacy with fairness [12].

Interactions between OGF's layers are bidirectional and recursive (see figure 2), Tier 1 (T1) sends consent logs and federated learning gradients to Tier 3 (T3), where DAECs audit system compliance and adjust privacy budgets (ϵ, δ)

as needed [12], meanwhile, Tier 2 (T2) uses edge AI to preprocess data locally, helping to reduce latency. To promote fairness in federated training, synthetic data oracles feed curated datasets into T1. At the same time, neuromorphic auditors in T3 monitor T2's outputs for signs of bias drift, triggering DAEC votes when model constraints need to be updated. If issues are found, self-healing mechanisms can automatically retrain models using data from T1 [17]. To support accountability, compensatory tokens (issued through T1's blockchain) are used to validate redress claims, closing the loop between governance and system execution [10].

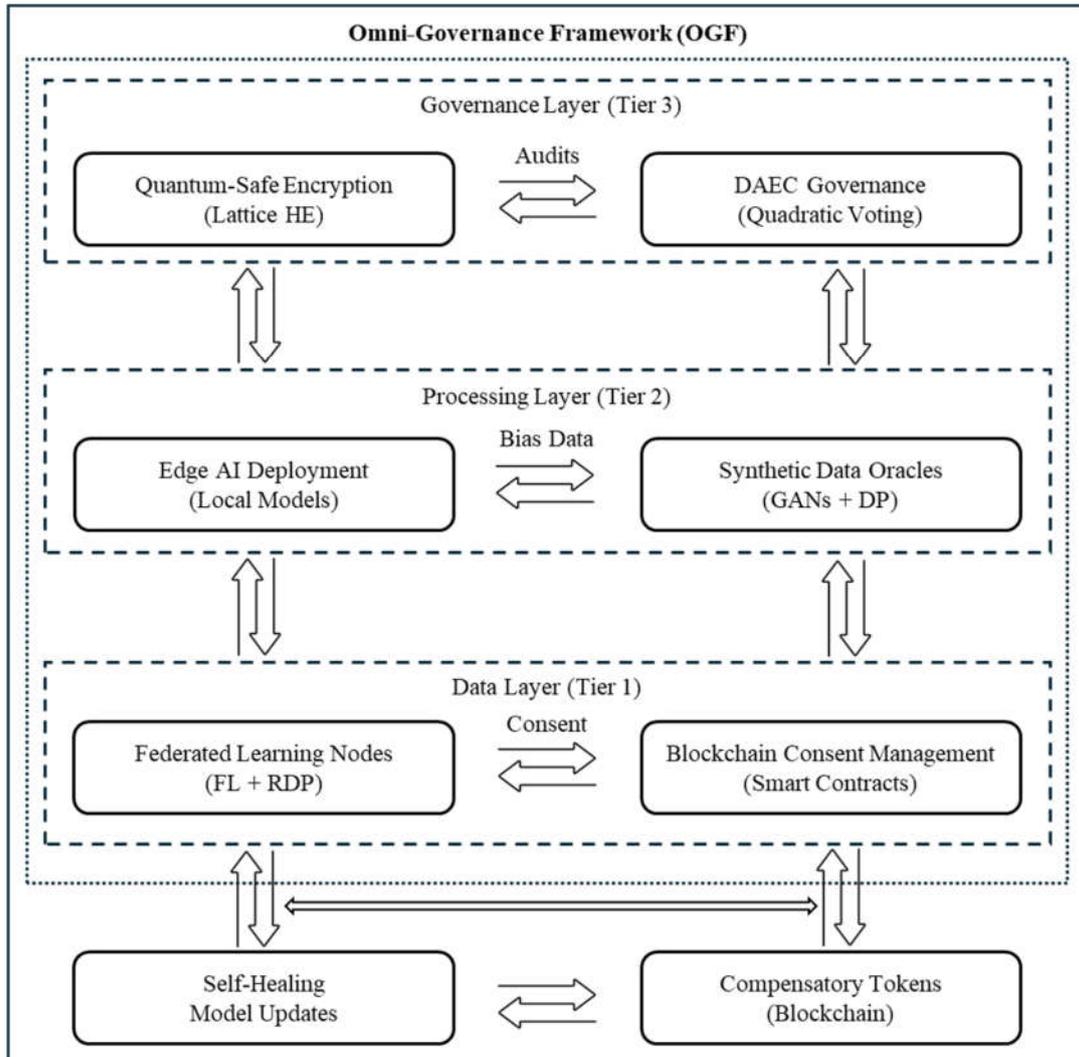


Figure 1: Architecture of Omni-Governance Framework (OGF)

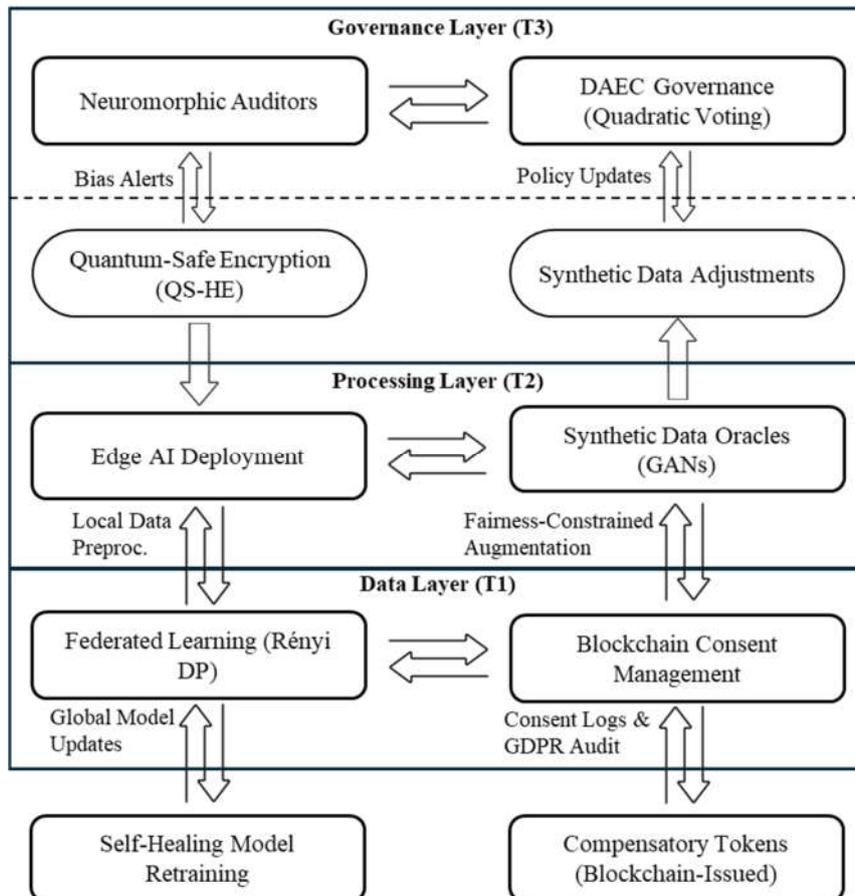


Figure 2: Workflow of Omni-Governance Framework (OGF)

4. Discussion

On the other hand, The Omni-Governance Framework (OGF) distinguishes itself from other existing federal learning and privacy-preserving frameworks, such as FATE (Federal AI Technology Enabled) and PySyft, because of its comprehensive integration of privacy, fairness, and governance mechanisms, specifically designed for Web 4.0 ecosystems.

While FATE and PySyft were pioneers in the fields of federated learning and secure multi-party computation, but they lack key features required for ethical AI in decentralized environments. For example, FATE's use of centralized differential privacy can actually worsen bias in imbalanced datasets, since the added noise tends to disproportionately affect underrepresented groups [24].

Although PySyft supports privacy-preserving federated learning but doesn't offer any built-in tools to address algorithmic fairness or meet regulatory requirements, leaving it to developers to patch together their own solutions.

OGF addresses these gaps by unifying Rényi differential privacy (RDP) with fairness-constrained synthetic data generation and mathematically linking privacy budgets (ϵ) to fairness metrics ($\Delta F \Delta F$) [27]. Unlike FATE's fixed approach to differential privacy DP, Rényi differential privacy (RDP) provides tighter privacy loss bounds, enabling better fairness-utility tradeoffs in decentralized settings [31]. Additionally, OGF's Decentralized Autonomous Ethics Committees (DAECs) resolve governance challenges that both FATE and PySyft ignore. Decentralized Autonomous Ethics Committees (DAECs) use quadratic voting to democratize decisions regarding privacy and fairness boundaries to ensure that stakeholders (not central coordinators) control ethical trade-offs [19]. This contrasts sharply with FATE's reliance on a central server for model aggregation and PySyft's lack of governance protocols.

Another key difference lies in quantum readiness, where both FATE and PySyft depend on classical encryption like AES-256, which is expected to be vulnerable to quantum attacks, especially through algorithms like Shor's, leaving these systems exposed in a post-quantum world [31].

OGF integrates lattice-based homomorphic encryption (e.g., CRYSTALS-Kyber) to secure federated gradients against quantum decryption, a feature absent in both FATE and PySyft [24]. Furthermore, OGF's neuromorphic auditors, spiking neural networks (SNNs) deployed on low-power edge devices, detect bias in real time, reducing energy costs by 100x compared to GPU-dependent fairness tools in FATE [18].

Regulatory compliance further differentiates OGF. While PySyft supports GDPR-aligned federated learning, it lacks automated mechanisms for consent revocation or bias redress [29]. OGF's blockchain-based consent management automates GDPR's "right to erasure" via zero-knowledge proofs, and its compensatory token system issues blockchain-backed redress for biased outcomes [5]. Neither FATE nor PySyft offer comparable accountability features (see table 2).

Table 2: Summary comparison between OGF, FATE and PySyft

Feature	OGF	FATE	PySyft
Privacy	Rényi DP + Quantum-Safe Encryption	Centralized Differential Privacy	Federated Learning + DP
Fairness	Neuromorphic Auditors + Fair Synthetic Data	Basic Statistical Audits	Not supported
Governance	DAECs with Quadratic Voting	Not available	Not supported
Regulatory Compliance	Full GDPR + EU AI Act Alignment	Not compliant	Partial GDPR via Federated Learning
Decentralization	Blockchain + Edge AI	Centralized Coordinator	Federated Nodes Only
Quantum Readiness	Lattice-Based Homomorphic Encryption	Not supported	Not supported
Redress Mechanisms	Token-Based Redress + Self-Healing Models	Not available	Not available
Energy Efficiency	Neuromorphic Hardware (SNNs)	GPU-Dependent	CPU/GPU Hybrid

5. Conclusions

The Omni-Governance Framework (OGF) addresses the growing needs for holistic governance in decentralized AI systems in Web 4.0, where privacy, fairness, and regulatory compliance increasingly overlap. By integrating Rényi differential privacy (RDP), fairness-constrained synthetic data, and quantum-resistant encryption into a tiered architecture, OGF resolves the privacy-fairness tradeoff that plagues federated learning. Unlike siloed tools like FATE and PySyft, OGF creates meaningful cross-tier synergies, for example, using synthetic data to strengthen federated datasets and deploying neuromorphic auditors to detect bias in real time to ensure compliance with evolving regulations like the EU AI Act while maintaining technical scalability.

A key innovation of OGF lies in its mathematical rigor, where Rényi Differential Privacy (RDP) provides tighter privacy guarantees than standard differential privacy, allowing for measurable tradeoffs between fairness and utility, which defined by the relationship $\Delta F = \frac{\partial F}{\partial \epsilon} \cdot \Delta \epsilon$, meanwhile, lattice-based homomorphic encryption secures federated gradients against quantum attacks. Using quadratic voting, the DAEC governance model democratizes ethical decision-making, whereas expert oversight and built-in safeguards, such as Sybil resistance, help mitigate token concentration and ensure balanced participation.

Despite its strengths, OGF faces limitations. The use of quantum-safe encryption can cause significant computational overhead, potentially slowing down Tier 3 processes and requiring hybrid pipelines that balance classical and quantum approaches.

While Tokenized redress mechanisms offer an effective means to address algorithmic harm, they also risk commoditizing these harms, causing an ethical issue that require closer scrutiny. Furthermore, the OGF framework relies on edge AI and assumes a widespread IoT infrastructure, potentially excluding regions with limited technological resources.

Future research should move in three key directions: (1) optimizing lattice-based encryption for real-time federated learning, (2) integrating OGF with ethical frameworks designed for AGI to anticipate emerging risks, and (3) creating interoperability standards that support compliance across different legal jurisdictions. To achieve this, policymakers

and technology experts must work together to improve adaptive governance models that strike the right balance between decentralization and accountability.

In conclusion, the Omni-Governance Framework (OGF) marks a major step toward AI governance. By combining technical innovation with ethical foresight, its modular, tiered architecture helps organizations navigate the complex realities of Web 4.0, while maintaining user trust, and as decentralized systems become central to sectors like healthcare and finance, frameworks like OGF will be indispensable in ensuring that technological progress aligns with societal values, where privacy is non-negotiable, biases are preempted, and governance is democratized.

References

- [1] Teixeira, J. E., & Tavares-Lehmann, A. T. C. (2022). Industry 4.0 in the European union: Policies and national strategies. *Technological Forecasting and Social Change*, 180, 121664.
- [2] Chu, W. (2022, May). A decentralized approach towards responsible AI in social ecosystems. In *Proceedings of the International AAAI Conference on Web and Social Media* (Vol. 16, pp. 79-89).
- [3] Truong, N. B., Sun, K., Lee, G. M., & Guo, Y. (2019). GDPR-compliant personal data management: A blockchain-based solution. *IEEE Transactions on Information Forensics and Security*, 15, 1746-1761.
- [4] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2), 1-19.
- [5] Haque, A. B., Islam, A. N., Hyrynsalmi, S., Naqvi, B., & Smolander, K. (2021). GDPR compliant blockchains—a systematic literature review. *Ieee Access*, 9, 50593-50606.
- [6] Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., ... & Zhang, Y. (2019). AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5), 4-1.
- [7] Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency* (pp. 77-91). PMLR.
- [8] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6), 1-35.
- [9] Pehlivan, C. N. (2024). Report: The EU Artificial Intelligence (AI) Act: An Introduction. *Global Privacy Law Review*, 5(1).
- [10] Akther, A., Arobee, A., Adnan, A. A., Auyon, O., Islam, A. S. M., & Akter, F. (2025). Blockchain As a Platform For Artificial Intelligence (AI) Transparency. *arXiv preprint arXiv:2503.08699*.
- [11] Voss, W. G. (2019). Cross-border data flows, the GDPR, and data governance. *Wash. Int'l LJ*, 29, 485.
- [12] Gharavi, H., Granjal, J., & Monteiro, E. (2025). PQBFL: A Post-Quantum Blockchain-based Protocol for Federated Learning. *arXiv preprint arXiv:2502.14464*.
- [13] Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- [14] Tiwald, P., Ebert, A., & Soukup, D. T. (2021). Representative & fair synthetic data. *arXiv preprint arXiv:2104.03007*.
- [15] Bühler, M. M., Calzada, I., Cane, I., Jelinek, T., Kapoor, A., Mannan, M., ... & Zhu, J. (2023). Unlocking the power of digital commons: Data cooperatives as a pathway for data sovereign, innovative and equitable digital communities. *Digital*, 3(3), 146-171.
- [16] Behera, M. R., Upadhyay, S., Shetty, S., Priyadarshini, S., Patel, P., & Lee, K. F. (2022). FedSyn: Synthetic data generation using federated learning. *arXiv preprint arXiv:2203.05931*.
- [17] Dilmaghani, M. S., Shariff, W., Farooq, M. A., Lemley, J., & Corcoran, P. (2024). Optimization of event camera bias settings for a neuromorphic driver monitoring system. *IEEE Access*.
- [18] Agrawal, S. (2024). Harnessing Quantum Cryptography and Artificial intelligence for next-gen payment Security: A Comprehensive analysis of threats and countermeasures in distributed ledger environments. *International Journal of Science and Research*, 13(3), 682-687.
- [19] Alsagheer, D., Xu, L., & Shi, W. (2023). Decentralized machine learning governance: Overview, opportunities, and challenges. *IEEE Access*, 11, 96718-96732.
- [20] Lahariya, C., Sundararaman, T., Ved, R. R., Adithyan, G. S., De Graeve, H., Jhalani, M., & Bekedam, H. (2020). What makes primary healthcare facilities functional, and increases the utilization? Learnings from 12 case studies. *Journal of family medicine and primary care*, 9(2), 539-546.

- [21] Prajapati, C. (2025). Decentralized Finance (DeFi) and Cryptocurrencies: The Latest Thinking of People Towards the Blockchain and FinTech Industry (Doctoral dissertation, University of the Cumberland).
- [22] Rayhan, S. (2023). Ethical implications of creating AGI: impact on human society, privacy, and power dynamics. *Artificial Intelligence Review*, 11(2), 44-59.
- [23] Gardner, A. (2022). Responsibility, recourse, and redress: A focus on the three R's of AI ethics. *IEEE Technology and Society Magazine*, 41(2), 84-89.
- [24] Almatarneh, R., Aljaidi, M., Alsarhan, A., Alshammari, S. A., & Alshammari, N. H. (2025). KTCGM: Towards A novel solution for enhancing Kerberos-5 with threshold cryptography and ML-based anomaly detection. *International Journal of Innovative Research and Scientific Studies*, 8(3), 3646–3662. <https://doi.org/10.53894/ijirss.v8i3.7328>
- [25] Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2023). Auditing Large Language Models: A Three-layered Approach", *AI and Ethics*.
- [26] Tu, T., Azizi, S., Driess, D., Schaekermann, M., Amin, M., Chang, P. C., ... & Natarajan, V. (2024). Towards generalist biomedical AI. *Nejm Ai*, 1(3), A1oa2300138.
- [27] Joseph, D., Misoczki, R., Manzano, M., Tricot, J., Pinuaga, F. D., Lacombe, O., ... & Hansen, R. (2022). Transitioning organizations to post-quantum cryptography. *Nature*, 605(7909), 237-243.
- [28] Liu, Y., Fan, T., Chen, T., Xu, Q., & Yang, Q. (2021). Fate: An industrial grade platform for collaborative learning with data protection. *Journal of Machine Learning Research*, 22(226), 1-6.
- [29] Almatarneh, R., Aljaidi, M., Alsarhan, A., Alshammari, S. A., Alhamazani, F., & Alshammari, A. B., (2025). An integrated AI-blockchain framework for securing web applications, mitigating SQL injection, model poisoning, and IoT spoofing attacks. *International Journal of Innovative Research and Scientific Studies*, 8(3), 2759–2773. <https://doi.org/10.53894/ijirss.v8i3.7077>
- [30] Mironov, I. (2017, August). Rényi differential privacy. In 2017 IEEE 30th computer security foundations symposium (CSF) (pp. 263-275). IEEE.
- [31] Almousa, Mohammad & Al-Zou'bi, Samah & Askar, Sami & AlQawasmi, Khaled & Al-Sherideh, Alaa & Samara, Ghassan & Almatarneh, Rami & Khouj, Mohammed & Odeh, Mahmoud. (2024). IoT Security Based On Lightweight Cryptographic (LWC) Algorithms: A survey. 1-9. 10.1109/ACIT62805.2024.10876942.