

ODIN AI : Federated Machine Learning on Blockchain

Democratizing AI through Decentralisation of Data, Computation, and Models

ODIN AI Team

Abstract—The rapid pace of AI development has highlighted significant challenges in its creation and deployment, primarily due to the centralized control maintained by a few large corporations. Such an approach exacerbates biases within AI models due to a lack of effective governance and oversight. Furthermore, it diminishes public engagement and raises serious data protection concerns. The resulting monopolistic control over data and model outputs also poses a threat to innovation and equitable data usage, as users unknowingly contribute to data sets that serve the interests of these corporations.

ODIN AI democratizes AI development and alignment through on-chain incentive mechanisms. By promoting open-source development and data ownership, ODIN AI facilitates an open and collaborative environment where participants can contribute models, data, and computing resources with rewards determined by on-chain consensus. This approach improves transparency and collaboration at scale without introducing biases from centralized entities. Ultimately, ODIN AI enables diverse communities to develop purpose-built AI models, offering bespoke solutions tailored to their specific needs, revolutionizing the landscape of AI development and deployment.

I. INTRODUCTION

Spanning all fields, collaboration has historically catalyzed innovation. This is manifest in the case of the scientific and the digital. By pooling collective expertise, we have forged disruptive solutions at speed. At present, this ideal faces barriers when applied to AI development and deployment: notably, diminished public engagement, pervasive concerns regarding concentrated control, and data protection exerted by a handful of corporations. Meanwhile, blockchain technology [1], [2] has demonstrated its efficacy in multiple areas needing distributed corporations, such as decentralized finance [3], voting and governance. Research into and deployment of blockchain to transform AI development is now underway.

ODIN AI, predicated on community involvement and a staunch commitment to data protection, is poised to spearhead the democratization of AI ecosystem by using blockchain.

A. The Problems with Centralized Control over AI Creation

In the present day, the primary obstacle to innovation in the realm of AI is its centralized control. This centralized structure mandates that all AI training, decision-making processes, and data storage are controlled within a single entity or location [4]. This results in the following pitfalls:

Single Point of Failure: Vulnerability to disruptions from technical issues and cyberattacks.

Value Plurality: Lack of value plurality means biases of single entities are reflected in AI. With centralized institutions exerting absolute control over models [5], [6], the

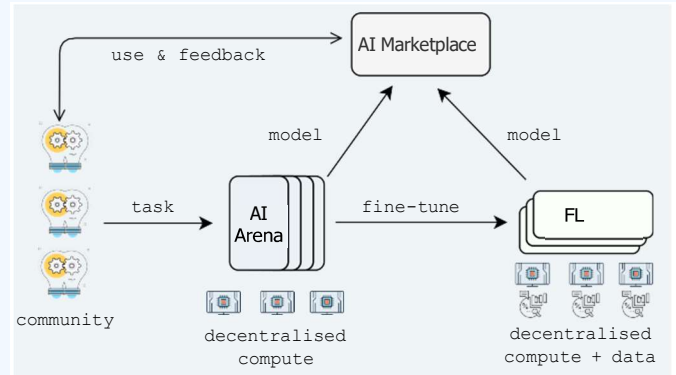


Fig. 1: ODIN AI System Logic. Upon a task creation, the model is first trained and validated in AI Arena, a blockchain-based decentralized training platform, and then optionally further fine-tuned in FL Alliance using participants' local data. Finally, the model is deployed by applications in the AI Marketplace, where feedback will be used to further improve the model.

values of the output models are also centralized [7]. For instance, the world reimaged by Google's generative AI tool, Gemini, is widely criticized [8].

Data Protection: Providers of closed-source Large Language Models (LLMs) [9], such as OpenAI, have the capability to monitor all user interactions with their models, thereby raising significant data protection concerns. In addition, under this centralized framework, every user who interacts with a LLM becomes an unwitting contributor of data to these vast corporations that maintain ownership of the models. There is a pressing need to enhance the fairness of contribution incentives and to more accurately assess the value of user-contributed data.

Governance: Recent research [10], [11], [12] has highlighted a concerning trend in which the lack of governance has led to a pronounced exacerbation of biases and inaccuracies within the models.

Scalability: As the volume of data and complexity of tasks increase, limited processing power acts as a bottleneck.

Innovation: Progress is stifled in an environment where a limited number of entities have the means to experiment.

B. ODIN AI's Solution

ODIN AI [13], [14] is a blockchain-based platform for decentralized AI. As shown in Figure 1 and 2, ODIN AI eliminates obstacles that prevent active participation in AI systems, em-

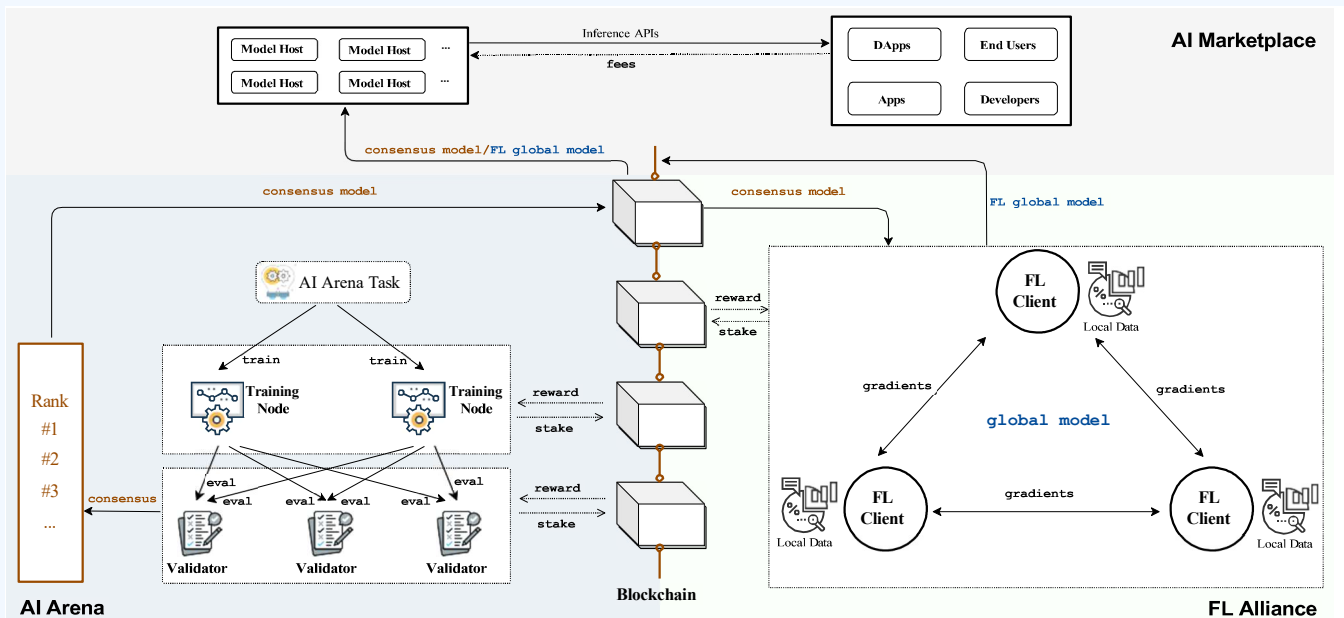


Fig. 2: ODIN AI System Overview. When a task is created in AI Arena, it is first trained by training nodes. These nodes then submit their models to validators, who evaluate and propose scores for each submission. The validators reach a consensus on these scores to determine the ranking of the submitted models. The consensus model can then be assigned to FL clients, who fine-tune and improve it using their local data, resulting in the FL global model. The AI Arena consensus model or the FL global model can be deployed and hosted in the AI Marketplace, providing interfaces to various applications. AI Arena train nodes, validators and FL clients need to stake to participate the system, and will be rewarded based on their performance.

powering communities to contribute models, data, or computing resources in a modular and decentralized way. AI models can be trained and validated in AI Arena and further refined in Federated Learning (FL) Alliance. Harnessing blockchain technology, ODIN AI introduces incentive mechanisms for participants, fostering a collaborative environment. This results in the development of a wide range of purpose-built models, created by, with, and for the communities, offering tailored solutions to meet specific needs.

II. ODIN AI SYSTEM OVERVIEW

The ODIN AI system consists of the blockchain layer, AI layer, and various participants. Each component plays an essential role in ensuring the system’s functionality and security.

A. Blockchain Layer

ODIN AI’s tokenomics incorporates a blockchain-based reward mechanism designed to enhance resilience against malicious user attacks. This robust security feature is underpinned by a carefully designed incentive mechanism. The blockchain layer acts as the foundation for both stakeholder participation and the distribution of rewards. This layer employs smart contracts to ensure that participants can securely lock in their stakes, fostering an environment of trust and transparency. The process is designed to incentivize participation by allocating rewards based on contributions, thus encouraging a more engaged and active community.

The blockchain layer’s inherent security features safeguard against fraudulent activities, ensuring uncompromised integrity

of staking and reward distribution. It is a critical component to support the model safety and improve resilience against malicious user attacks. By leveraging smart contracts, the system automates an efficient and fair rewards process. Automation reduces human error and ensures that rewards are distributed in a timely and fair manner.

B. AI Layer

The AI layer offers infrastructure for decentralized training, extracting and monetizing knowledge from data. It encourages compute and data contributions from the community, using blockchain for rewards based on their contributions.

- **AI Arena.** AI layer supports a conventional machine learning (ML) model training paradigm, optimizing models directly on users’ devices with their own or public data. To maximize the generalization ability and performance of the final trained models, this layer is designed to encourage community members to contribute various public or local data, harnessing the broader community’s power. By leveraging blockchain, it ensures contributors are continually engaged and rewarded based on the quantifiable impact of their data on improving the models.
- **FL Alliance.** Utilizing the FL [15] approach, the AI layer enables thousands of participants to collaboratively train a global model, where data sovereignty is preserved by ensuring that no local data are transmitted at any stage of the training process. Within the AI layer, a model aggregation component allows participants to upload weights

from models trained on their unique local data. These weights are then aggregated to build an optimal global model, enhancing its generalization capabilities and performance. The integration of training task automation and deployment orchestration components simplifies the process for users to join tasks and contribute valuable knowledge extracted from their data.

In ODIN AI, AI Arena tasks will engage participants from the Web2 AI community, who possess the necessary computational resources to train and validate models using publicly available datasets. These trained models can be further refined through FL Alliance tasks, which draw in participants capable of contributing their own local data.

C. AI Marketplace

Once models are trained and fine-tuned through AI Arena and FL Alliance, they can be hosted on our platform. This platform serves as a comprehensive environment for deploying ML models, making them accessible within blockchain networks of virtual machines (VMs). By integrating with these networks, the platform facilitates the seamless execution and inference of complex ML models, providing real-time, scalable, and secure solutions.

The infrastructure for model management includes version control, model monitoring, and automated updates, ensuring that the models remain accurate and efficient over time. It can provide inference APIs or SDKs that developers can use to integrate these models into their applications.

Model hosts are compensated based on the quality and frequency of their contributions. They play a crucial role in generating inferences and maintaining the infrastructure.

D. Participants

There are various categories of participants in ODIN AI.

- 1) **Task Creators:** Task creators will define the training tasks. Any participant who is willing to stake sufficient assets into the system or has already contributed to the system can potentially be selected as a task creator. This broadens the range of stakeholders, conferring a sense of ownership and active involvement.
- 2) **Training Nodes:** Training nodes compete in AI task training and are required to stake tokens to be eligible. This requirement ensures a commitment to the network's integrity and facilitates a distributed, trust-based mechanism for task assignment. This stake acts both as a gatekeeper to maintain a high standard and as a foundational element in the network's security protocol, ensuring that nodes have a vested interest in proper execution and the overall health of the ecosystem.
- 3) **Validators:** Validators are responsible for evaluating work done by training nodes, submitting validation scores that influence reward distribution. They participate by staking tokens, which grants them the opportunity to validate tasks assigned to them, ensuring hardware compatibility and fair task distribution proportional to their stake. Upon completion of a task, they can withdraw

their stake and claim rewards, which are calculated based on their performance and adherence to the expected outcomes. The design ensures that validators are incentivised to provide accurate and honest validations, thereby maintaining the quality and reliability of the network's computational tasks.

- 4) **Delegators:** Delegators contribute to the ODIN AI system by supporting other participants' staking process, enhancing the network's validation capacity without directly participating in the task training or validation process. Delegators share in the rewards earned by their associated delegates, based on predefined algorithms that account for their staked contribution. Note that training nodes and validators who choose to accept delegation are free to choose a reward share ratio. The higher the ratio, the bigger the reward share their delegators will receive. The role of delegators allows individuals to participate in the network's training, validation and economic activities, leveraging their tokens to support delegates, without needing the technical capabilities to train or validate tasks themselves.
- 5) **FL Clients:** With a FL framework, FL clients will contribute their local data to enhance the model trained for the AI Arena task. In each FL task, participants will be randomly designated as either proposers or voters. Proposers will be tasked with training the model within a FL framework, while voters will assess the training outcomes produced by the proposers. Both proposers and voters will receive rewards or face penalties based on their respective performances. ODIN AI ensures that all participants are motivated to contribute effectively to the overall model improvement.
- 6) **Model Hosts:** The role of a model host in AI Marketplace involves deploying and managing trained models, providing infrastructure for secure and scalable execution, and enabling access through APIs and SDKs. The host ensures the models are kept up-to-date, monitors their performance, and facilitates integration into applications. Additionally, they will be compensated for their contributions to generating inferences and maintaining the system's integrity.

III. ODIN AI TOKENOMICS

ODIN AI aims to build a fair and incentive-compatible ecosystem, designed to foster collaboration and ensure long-term alignment within its community. This vision is realized through a strategically designed reward allocation system, an effective slashing mechanism for accountability, and the cultivation of active token demand.

A. Token Supply

1) **Emission:** ODIN AI's ecosystem will feature \$ODIN AI tokens, set to be distributed to various stakeholders through an initial token emission and a strategically designed reward allocation system over time. Participants will receive rewards in \$ODIN AI tokens based on their contributions to the system.

Participants in the ODIN AI system, such as training nodes and validators, are required to contribute computing or storage resources to complete model training and validation in order to receive rewards. This means that the value of the $\$ODIN AI$ token will, at a minimum, correspond to the value of the resources consumed during these processes.

2) *Slash*: ODIN AI robust mechanisms ensure the integrity and reliability of the system by penalizing participants that engage in malicious activities. In AI Arena, if a participant is identified as acting against the system’s rules or attempting to undermine the system through malicious actions, they are subjected to “slashing”. Slashed tokens will be rewarded to the honest participants. Slashing protects the system from immediate threats by disincentivizing malicious actors and reinforces a culture of trust and cooperation among participants.

B. Token Demand

Active token demand is encouraged through multifaceted approaches as follows, showing the value of circulating tokens within the ecosystem.

1) *Utility*: Participants are required to stake $\$ODIN AI$ to play a role. This reflects their vested interest in the integrity and success of operations. For task creators facing urgent needs to gather top-notch trainers for their model training or operating under tight deadlines, they may opt to pay additional $\$ODIN AI$ as bounties. These bounties will then be distributed as payments to participants involved in those specific tasks to prioritize the training processes. Participants can be also supported by delegators through $\$ODIN AI$ token delegation. By doing so, the system boosts the participants’ stake within the ODIN AI system and incentivizes a symbiotic relationship. Delegators, in turn, earn a share of the rewards earned by their participants, fostering a competitive environment where participants are motivated to offer attractive terms to potential delegators.

2) *Payment*: Community members are able to access and utilize winning models which are trained and fine-tuned in AI Arena and FL Alliance, and hosted on AI Marketplace. End users enjoy rate limit in their access to such models based on their stake amount, beyond which they will be charged in $\$ODIN AI$ as payment. On the other hand, model hosts need to stake $\$ODIN AI$ in order to host winning models. They are able to customize whether and how to charge end users of these models. At inception phase, model hosts will receive part of the daily emission in order to incentivize their participation. Yet such incentives are expected to diminish over time. Overall, such design creates a sustainable and competitive environment in which demand and supply for cutting-edge models are dynamically balanced, fostering innovation and ensuring that the latest advancements continue to meet the evolving needs of the market. The payment mechanism also creates a non-negligible financial barrier in access to our models, thus helps mitigate potential DoS attacks from malicious participants.

3) *Governance Participation*: Holding $\$ODIN AI$ tokens grants members the power to influence the network’s fu-

ture through participation in the Decentralised Autonomous Organization (DAO) governance. This not only decentralizes decision-making but also adds a layer of utility and value to the tokens, as they become a key to shaping the ecosystem’s development.

IV. ODIN AI INCENTIVE AND SECURITY

A. Incentive

ODIN AI leverages well-designed incentive mechanisms to reward participants. The distribution of newly emitted tokens is carefully orchestrated across AI Arena tasks and FL Alliance tasks, reflecting a strategic allocation that hinges on the staking dynamics within each task category.

In our system, verified tasks are granted a share of daily rewards, serving as an incentive to foster the growth of the task creation ecosystem. This reward distribution is intentionally restricted to tasks approved by the DAO to safeguard the protocol from being exploited by low-quality or malicious tasks that could otherwise drain emissions without contributing meaningful value.

Each newly created AI Arena and FL Alliance task has the option to undergo a verification process conducted by the community-led DAO. This process is designed to ensure that tasks meet the necessary standards of quality and alignment with the ecosystem’s goals. Once a task successfully passes verification, it becomes eligible for $\$ODIN AI$ ’s daily emissions, providing the task creator with additional resources to incentivize participation and collaboration.

On the other hand, if a task is created permissionless without the ODIN AI DAO’s verification, the responsibility falls on the task creator to self-fund the task. This involves using their own $\$ODIN AI$ to cover the costs associated with reward allocations for various participants. While this route allows for greater flexibility and decentralization in task creation, it also places the financial burden of supporting the task’s ecosystem on the creator. This mechanism is designed to balance innovation with quality control, ensuring that only well-constructed tasks benefit from community-supported rewards while still allowing for creative freedom in the ecosystem.

In the long run, this dual approach aims to encourage high-quality task creation, foster a vibrant and trustworthy ecosystem, and maintain the integrity of the $\$ODIN AI$ reward system by aligning incentives with the community’s standards and goals.

Once tasks are created, the distribution of rewards between DAO-verified AI Arena and FL Alliance tasks is dependent on their relative stake number of active tasks. As such, the rewards of $\$ODIN AI$ allocated to all active AI Arena tasks will be:

$$R^{AI} = C_0 \cdot \frac{S^{AI}}{S^{AI} + S^{FL}}$$

and for all active FL Alliance tasks:

$$R^{FL} = C_0 \cdot \frac{S^{FL}}{S^{AI} + S^{FL}}$$

Attacks	Description	ODIN AI Mitigation
Sybil Attacks	An attacker might gain disproportionate influence in the ODIN AI system by creating and controlling multiple fake identities of participants.	<ul style="list-style-type: none"> ► Staking assets increases the difficulty of controlling many training nodes or validators. ► Blind validation mechanism prevents collusion between training nodes and validators. ► In each task, only the top k_1 training nodes and the top k_2 validators will be rewarded, ensuring that participants with poor performance do not receive rewards.
DoS Attacks	An attacker might exhaust the ODIN AI system resource and make it unavailable to honest participants.	<ul style="list-style-type: none"> ► Rate limiting is implemented to restrict the frequency and volume of actions within a certain time frame, ensuring that no single participant can overwhelm the system.
Free-rider Attacks	Free riders benefit from a system without contributing fairly. In the ODIN AI system, a free rider training node may randomly submit models without actual training. Similarly, free rider validators give random scores instead of honestly evaluating models.	<ul style="list-style-type: none"> ► In each task, only the top k_1 training nodes and the top k_2 validators will be rewarded, ensuring that participants with poor performance do not receive compensation. ► ODIN AI Arena consensus guarantees that honest participants who contribute diligently are appropriately recognized and rewarded, deterring free riders from exploiting the process.
Lookup Attacks	Training nodes could cheat by learning to predict past validation score calculations.	<ul style="list-style-type: none"> ► Two datasets, i.e., Datasets A and B, are used as validation sets to evaluate the models. Consequently, even if a training node manages to optimize its performance for Dataset A, it could still underperform on Dataset B. By carefully calibrating the rewards between Dataset A and B, ODIN AI effectively motivates training nodes towards developing genuinely high-quality models.
FL Model Poisoning Attacks	In FL Alliance, an attacker may use biased or corrupted data during the training process to degrade the model's performance.	<ul style="list-style-type: none"> ► By aggregating contributions, majority voting minimizes the impact of single malicious participants. ► The slashing mechanism penalizes malicious clients, deterring model poisoning by reducing their rewards and discouraging future attacks.

TABLE I: Summary of how ODIN AI system mitigation against potential attacks.

in which C_0 is the daily emission of $\$ODIN AI$, S^{AI} refers to the total stake amount of all active AI Arena tasks and S^F refers to the total stake amount of all active FL Alliance tasks. Note that at the initial phase; to incentivize participation, task creators will also receive a slice of the reward pool. This reward, however, is expected to be phased out over time.

In AI Arena, this allocation is meticulously calculated based on the aggregate staking contributions from task creators, training nodes, validators, and delegators for each task.

1) *Rewards among AI Arena Tasks*: Within the span of a single day, consider the situation where there are M AI Arena tasks with the total staking amounts of (S_1, \dots, S_M) . The total staking amount, S_i , includes the stakes from all participants involved in this task. This means that the stakes from any type of user will influence the reward distribution among tasks. p is a system parameter that can be adjusted via DAO decision.

Assume the amount of daily emitted $\$ODIN AI$ token is C_{AI} .

For an AI Arena task with the total staking amount of S_i , its daily total rewards is:

$$R_i^{AI} = C_{AI} \cdot \frac{S_i^p}{\sum_{k=1}^M S_k^p}$$

For each AI Arena task, rewards are allocated among task

creators, training nodes, validators, and delegators. In the initial version of ODIN AI, if a validator has delegators, 1% of their rewards are designated for these delegators. It is important to note that this distribution parameter is flexible and subject to adjustments through the ODIN AI DAO governance.

2) *Rewards among FL Alliance Tasks*: A FL Alliance task should be derived from a finished AI Arena task to be further fine-tuned. The initiation of an FL task automatically triggers the creation of a new FL contract. For each active FL Alliance task within the ecosystem, daily rewards are transferred to the respective FL smart contract, provided the task is still in progress and has not surpassed its maximum allotted lifecycle. This preliminary step ensures that the rewards are earmarked and protected for participants actively engaged in the task. Subsequently, upon meeting the predefined conditions, the FL smart contract autonomously distributes the rewards to the participants, according to their contributions.

B. Security

As shown in Table I, the ODIN AI system's security is de-

signed to be resilient against attacks.

Sybil attacks are mitigated by a requirement to stake a minimum amount of assets, making it costly to control

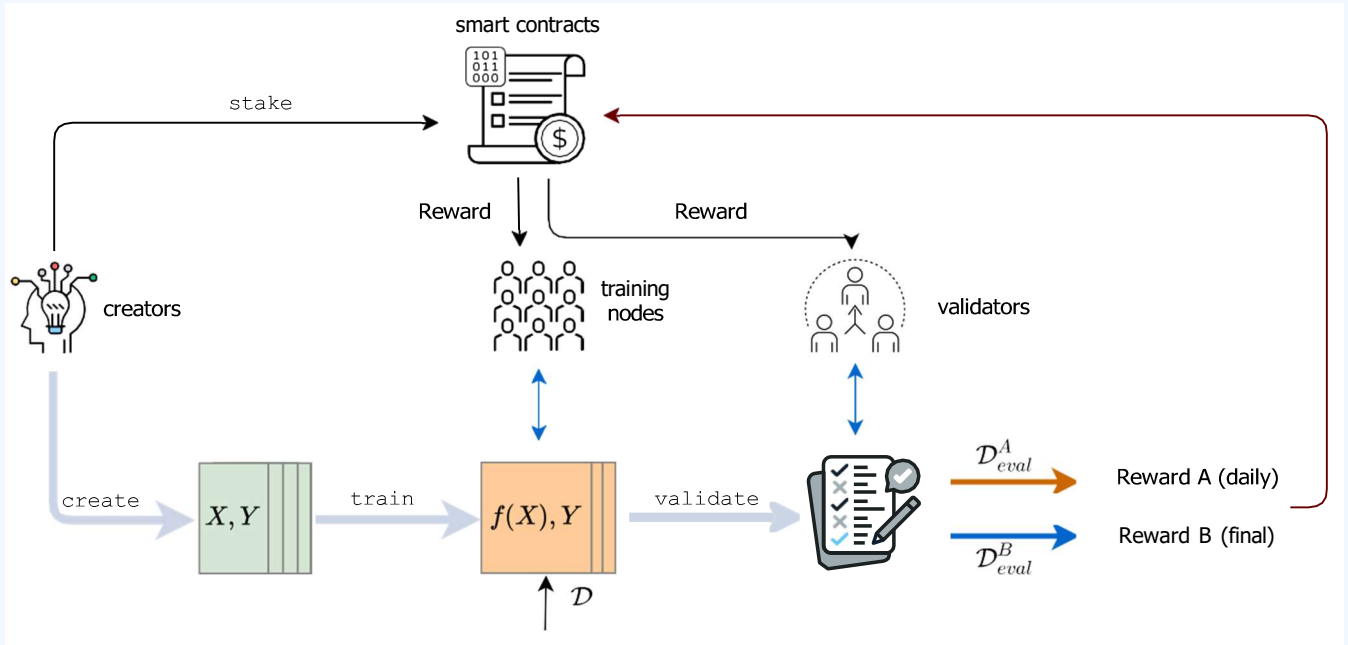


Fig. 3: Overview of the workflow of an ODIN AI Arena task. Validators earn rewards based on their consensus scores. Two types of rewards are used to incentivize training nodes in order to mitigate their lookup/overfitting attacks.

multiple identities. Validators are kept unaware of the model origins, reducing the risk of collusion. Only the top-performing training nodes and validators receive rewards, discouraging poor performance and manipulation. To mitigate DoS attacks, the system implements rate limiting, preventing any single participant from monopolizing resources. Free-rider attacks are addressed by rewarding only the top contributors, ensuring that participants who do not genuinely contribute cannot benefit. The use of dual datasets (Dataset A and B) in evaluations prevents lookup attacks, as optimizing for one dataset does not guarantee success in the other. For FL model poisoning attacks, a majority voting system and slashing mechanism protect the model’s integrity, punishing malicious actors and discouraging future attempts. These measures collectively fortify ODIN AI against a range of threats, promoting a secure and reliable decentralized training environment for participants.

V. ODIN AI CONSENSUS IN AI ARENA

Figure 3 shows the overview of the workflow of an ODIN AI Arena task.

A. Task Creators

Task creation is the primary stage of the training cycle. Creators define the desired models and submit tasks to the platform. Anyone who satisfies the criteria is eligible to be a task creator, making the system inherently democratic and accessible to a wide range of stakeholders. This inclusivity fosters a sense of ownership and active involvement within the ODIN AI community.

To qualify as a task creator, users must meet one or more of the following criteria:

- Stake a sufficient amount of $\$ODIN_{AI}$.
- Have successfully trained or validated a task previously, as evidenced by on-chain records.
- Possess a reputation in the ML space or be recognized as a domain expert in relevant fields, as verified by the ODIN AI community.

If the task creator and the created task are verified by the ODIN AI DAO, the task will be eligible for daily $\$ODIN_{AI}$ emissions. However, if the task creator chooses not to undergo verification by the community-led DAO, they must self-fund the task using $\$ODIN_{AI}$ to cover the costs associated with reward allocations for the participants.

In addition to gaining access to the desired training model, task creators may also earn rewards for their contributions. However, these rewards are expected to be gradually phased out over time.

B. Training Node and Validator Selection

In this setup, each participant first stakes in the system to be eligible to perform task training or validation.

In practice, rate limiting is adopted to determine the number of times participants can access validation for a given task. As illustrated in Figure 4, the likelihood of a participant being selected to validate a task submission increases with their stake. However, the rate at which validation frequency increases relative to the staking amount tends to diminish as the staking amount grows.

C. Training in AI Arena

We consider the dataset held by the training node, \mathcal{D}_{local} , which contains locally sourced data samples, comprising feature set X and label set Y , with each sample $x_i \in X$




Fig. 4: Example: rate limiting for validation frequency.

corresponding to a label $y_i \in Y$. We define a predictive model f , aiming to learn patterns within D such that $f(x_i) \approx y_i$.

To quantify the prediction metric, accuracy as an example, the task trainer will introduce a loss function $L(f(x_i), y_i)$, assessing the discrepancy between predictions $f(x_i)$ and actual labels y_i . A generic expression for this function is:

$$L = \frac{1}{N} \sum_{i=1}^N l(f(x_i), y_i)$$

where N denotes the total sample count, and l signifies a problem-specific loss function, e.g., mean squared error or cross-entropy loss.

The optimisation goal is to adjust the model parameters θ to minimise L , typically through algorithms such as gradient descent:

$$\theta_{\text{new}} = \theta_{\text{old}} - \eta \nabla_{\theta} L$$

where η represents the learning rate, and $\nabla_{\theta} L$ the gradient of L with respect to θ . Utilising the aggregated dataset D , parameter θ is iteratively updated to reduce L , consequently improving the model's predictive accuracy. This optimisation process is conducted over a predefined number of epochs E , each epoch consisting of a complete pass through the entire dataset D .

D. Validation in AI Arena

Consider a selected group of validators, denoted as $V_j \in V$, each equipped with the evaluation dataset D_{eval} from the task creator. This dataset consists of pairs (x_i, y_i) , where x_i represents the features of the i -th sample, and y_i is the corresponding true label.

The model, trained by designated training nodes, is denoted as θ_p^{task} . The primary objective of θ_p^{task} is to predict the label \hat{y}_i for each feature vector x_i contained within D_{eval} .

To assess the performance of θ_p^{task} on D_{eval} , we use an general evaluation metric denoted by eval . Here, we exemplify with accuracy, which is calculated as follows:

$$\text{eval}(\theta_p^{\text{task}}, D_{\text{eval}}) = \frac{1}{|D_{\text{eval}}|} \sum_{(x_i, y_i) \in D_{\text{eval}}} \mathbf{1}(y_i = \hat{y}_i)$$

Here, $\mathbf{1}$ represents the indicator function that returns 1 if the predicted label \hat{y}_i matches the true label y_i , and 0 otherwise. The function $|D_{\text{eval}}|$ denotes the total number of samples within the evaluation dataset.

Each predicted label \hat{y}_i from the model θ_p^{task} is compared against its corresponding true label y_i within the dataset D_{eval} . The calculated metric result (accuracy here) serves as a quantifiable measure of θ_p^{task} 's effectiveness at label prediction across the evaluation dataset.

E. Reward for Training Nodes in AI Arena

Within a single AI Arena task, the reward distribution between training nodes and validators is determined based on their relative stake amounts.

We assume there are n submissions (O_1, \dots, O_n) from n training nodes, and m validators (V_1, \dots, V_m) , each with stakes (s_1, \dots, s_m) . The stakes represent the validators' commitment and trust in the process, influencing the weight of their evaluations in the aggregated score.

Let the total daily reward allocated to a task be denoted as R_0 and the parameter γ controls the split rewards, defining the balance between fixed and stake-dependent reward components. Then, the total rewards for training nodes are:

$$R_0 \cdot \gamma + (1 - 2\gamma) \cdot \frac{\sum_{i=1}^n t_i}{\sum_{i=1}^n t_i + \sum_{j=1}^m s_j}$$

- Each validator $V_j (1 \leq j \leq m)$ evaluates the n models submitted by the training nodes, producing a score vector $r_j = (r_{j1}, \dots, r_{jn})$. These scores reflect the perceived accuracy, reliability, or performance of each model according to predefined criteria. The outlier scores proposed by malicious validators will be ignored by honest validators before taking into account in the following steps.
- The final score for each model from the training nodes is determined through a weighted aggregation:

$$r^{\rightarrow} = \left(\frac{\sum_j r_{j1} \cdot s_j}{\sum_j s_j}, \dots, \frac{\sum_j r_{jn} \cdot s_j}{\sum_j s_j} \right)$$

This means that the evaluations of validators with higher stakes have a larger impact on the final outcome.

- We then compute the following geometry series:

$$g_k = \frac{1 - q}{1 - q^m} \cdot q^{k-1}$$

in which k denotes a given training node's rank amongst its peers in the same task, whereas q represents the common ratio of the geometric series and m is the number of training nodes in a given task.

Participants in the ODIN AI system, such as training nodes and validators, are required to contribute computing or storage resources to complete model training and validation in order to receive rewards. This means that the value of the $\$ODIN$ AI token will, at a minimum, correspond to the value of the resources consumed during these processes.

- We finally compute the total rewards allocated for the training nodes as well as their delegators, which is based on the quality of their submission and their total amount of stake:

$$f_i(g_i, t_i) = \frac{g_i \cdot t_i^{\alpha}}{\sum_{k=1}^n g_k \cdot t_k^{\alpha}}$$

where α_t is a system parameter which determines the influence of the stake on the reward distribution, and t_i the total stake amount of from the training node i as well as its respective delegators.

- Consider σ as the reward ratio set by training node i itself which determines the ratio of rewards shared between training node i and its respective delegators. Consider also that a training node i ' stake in the task is t_n and stakes delegated to training node i is t_d , i.e. $t_i = t_n + t_d$, then the actual reward for training node i is:

$$f_i \cdot \sigma + (1 - \sigma) \cdot \frac{t_n}{t_n + t_d}$$

F. Reward for Validators in AI Arena

Similar to reward calculation for training nodes, the rewards for all validators in the same given AI Arena task is:

$$R_0 \cdot \gamma + (1 - 2\gamma) \cdot \frac{\sum_{j=1}^m S_j}{\sum_{i=1}^n t_i + \sum_{j=1}^m S_j}$$

For each validator V_j , we compute the distances between their score and the final aggregated score:

$$\begin{aligned} \vec{\Delta}_j &= (\Delta_{j1}, \dots, \Delta_{jn}) \\ &= \left(\frac{\sum_i r_{j1} \cdot S_j}{\sum_j S_j} - r_{j1}, \dots, \frac{\sum_i r_{jn} \cdot S_j}{\sum_j S_j} - r_{jn} \right) \end{aligned}$$

We define a distribution function f_i , which satisfies:

- $f_i(\Delta_{1i}, S_1) + \dots + f_i(\Delta_{mi}, S_m) = 1$,
- f_i decreases over the distance Δ_{ji} ,
- f_i increases over the stake amount S_j .

To fulfill the three criteria, we can employ a modified version of the Softmax Function:

$$f_i(\Delta_{ji}, S_j) = \frac{e^{-\lambda_v \Delta_{ji}} \cdot S_j^{\alpha_v}}{\sum_{k=1}^m e^{-\lambda_v \Delta_{ki}} \cdot S_k^{\alpha_v}}$$

The parameters λ and α play crucial roles:

- λ_v
 - Purpose: Controls the sensitivity of the function to the distance Δ_{ji} . This distance measures the discrepancy between a validator's score and the aggregated score.
 - Effect: A higher λ_v increases the function's sensitivity to score accuracy, emphasising the importance of precise evaluations.

- Selection Criteria: The choice of λ_v balances the need to penalise inaccuracies against the goal of rewarding nearly accurate evaluations.

• α_v

- Purpose: Determines the influence of the stake amount S_j on the reward distribution, thereby adjusting the weight given to validators' financial contributions.
- Effect: Allows for balancing between the importance of validators' financial commitment and their performance accuracy. A higher α_v gives more weight to the stake amount in the reward calculation.
- Selection Criteria: Reflects the system's philosophy regarding the stake's importance relative to score accuracy. An α_v of 0 means stake amounts are ignored, while a higher value increases their impact.

If the validator finishes multiple (i.e., N) validation tasks, then it reward ratio is:

$$\sum_{i=1}^N f_i(\Delta_{ji}, S_j)$$

If a validator's stake in the task is s_v , and S_j is its accumulative stake by considering the total delegation amount S_d on this validator, i.e., $S_j = s_v + S_d$, the actual reward ratio for this validator is:

$$\sum_i f_i(\Delta_{ji}, S_j) \cdot \sigma + (1 - \sigma) \cdot \frac{s_v}{s_v + S_d}$$

where σ is a system parameter which controls the rewards split between a given validator and its delegators.

G. Delegate Staking

Delegators may entrust their tokens to participants of their choosing to receive a passive investment income stream. The receivers can thus amplify their stake, influence, voting power, and rewards. These rewards are shared with the delegators, furthering cooperation. This extends participation to users who have tokens but lack the technical expertise to perform AI model training or validation.

Specifically, reward for the delegator depends on:

- The quality of the training nodes or validators selected for delegation.
- The amount of stake delegator has delegated.

Formally, reward for delegator who delegates to a training node can be calculated as:

$$f_i \cdot (1 - \sigma) \cdot \frac{t_d}{t_n + t_d}$$

whereas f_i refers to the total reward distributed to the training node i and delegator based on the quality of the training node's submission, t_d is the stake amount from this given delegator, t_n is the stake amount from training node i and σ is the reward share ratio pre-determined by training node i .

	Total	Training Nodes				Delegators	Validators		
		Training Node A		Training Node B		Delegator A	Val. A	Val. B	Val. C
		Reward A (10%)	Reward B (90%)	Reward A (10%)	Reward B (90%)	($\sigma = 0.4$)			
Day 1	309,157.68	4,937.1	44,433.9	5,053.9	45,485.1	8,712.6	73,997	87,834	38,703
Day 2	309,157.68	4,937.1	44,433.9	5,053.9	45,485.1	8,712.6	73,997	87,834	38,703
Day 3	309,157.68	4,937.1	44,433.9	5,053.9	45,485.1	8,712.6	73,997	87,834	38,703
Day 4	309,157.68	4,937.1	44,433.9	5,053.9	45,485.1	8,712.6	73,997	87,834	38,703
Day 5	309,157.68	4,937.1	44,433.9	5,053.9	45,485.1	8,712.6	73,997	87,834	38,703

TABLE II: Example: Updated reward distribution for one task (Days 1–5) with 2 training nodes (A,B), Delegator A ($\sigma = 0.4$), and 3 validators (A,B,C). $\delta = 0.1$ for Reward A (10%) vs. Reward B (90%).

Similarly, reward for delegator who delegates to a validator can be calculated as:

$$\sum_i f_i(\Delta_{ji}, s_j) \cdot (1 - \sigma) \cdot \frac{s_d}{s_v + s_d}$$

in which $\sum_i f_i(\Delta_{ji}, s_j)$ is the Softmax function for validator mentioned above, whereas s_d refers to the stake amount from a given delegator and s_v is the stake amount from the validator itself.

In the future, ODIN AI delegate staking has the option to be integrated with existing restaking platforms to attract users from a border blockchain community.

H. Various Validation Sets

To mitigate the lookup attacks from malicious training nodes, ODIN AI validators adopt diverse validation datasets. Specifically, for a AI Arena task spanning x days, the validation dataset used during the initial $x - 1$ days differs from that of the final day. These distinct validation datasets are associated with two types of rewards: **Reward A** for the initial period and **Reward B** for the final day. This strategic approach enhances security by varying the data against which training nodes are validated, thereby complicating any potential malicious attempts to exploit predictable validation scenarios.

- For each AI Arena within the ecosystem, the rewards mechanism for training nodes is thoughtfully designed to comprise two distinct parts: **Reward A** and **Reward B**.
 - Reward A** provides a daily contingent reinforcement schedule, incentivizing participant engagement with AI Arena tasks through immediate gratification. This continuous reward mechanism fosters sustained participation by providing consistent feedback and reinforcing contributions.
 - Reward B**, implements vesting, releasing tokens upon successful task completion within a predetermined lifespan. This incentivizes participants to both engage in tasks and ensure their timely and efficient completion. Vesting also functions as a quality control mechanism, promoting focused contributions aligned with project deadlines.

- Reward A of the AI Arena task is:

$$R_i^{A,A} = \delta \cdot R_i^A$$

- Reward B of the AI Arena task is:

$$R_i^{A,B} = (1 - \delta) \cdot R_i^A$$

δ is a configurable system parameter.

I. Example

We consider the rewards for the participants in task 1.

We assume that:

- Daily total rewards for all AI Arena tasks for a given day is 309,157.68.
- There are 2 nodes and 3 validators in this given task.
- Nodes A and B stake 3,000 and 3,500 \$ODIN AI respectively, while Validators A, B and C stake 3,000, 6,000 and 3,000 respectively. Node A also receives an additional 1,000 ODIN AI from its delegators, which brings the t_i (total stake including delegated stake) to be 4,000 for Node A.
- For simplicity, we assume γ to be 0 in this example.
- We further assume σ to be 0.4, and δ to be 0.1.
- We consider the scores for Nodes A, B, Validators A, B and C are 0.501435, 0.498565, 0.472768, 0.280226 and 0.247006 respectively.

First, for this given task on Day 1, total rewards for all training nodes are:

$$309157.68 \times \frac{6500}{6500 + 12000} = 108623.7$$

We can then compute the rewards for Node A and its delegators. Consider $\alpha_{t=1}$, rewards for Node A (together with delegators) are:

$$\frac{0.501435 \times 4000}{(0.501435 \times 4000) + (0.498565 \times 3500)} \times 108623.7 = 58084$$

Similarly, for Node B, its rewards together with its delegators are:

$$\frac{0.498565 \cdot 3500}{(0.501435 \cdot 4000) + (0.498565 \cdot 3500)} \times 108623.7 = 50539$$

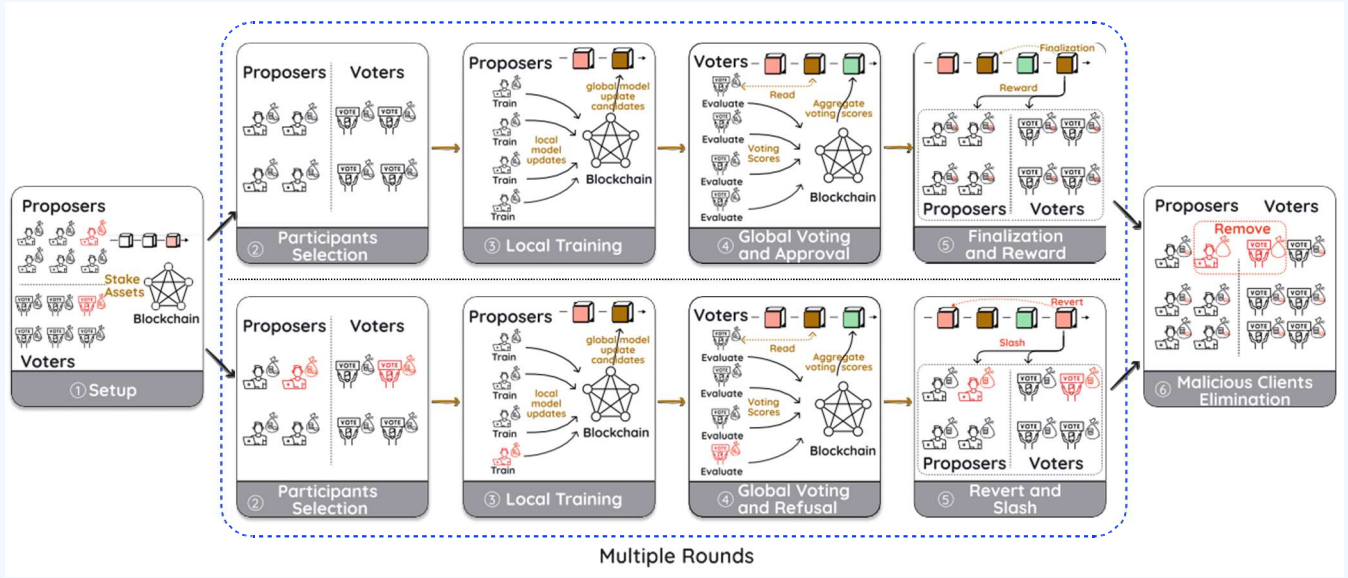


Fig. 5: Overview of the workflow of a ODIN AI FL task, adopted from our previous work [14]. During each round, FL clients are randomly selected to act as proposers or voters. Their staked amount can be either rewarded or slashed based on the outcome of the majority voting results. Malicious clients will ultimately be removed from the FL system.

Given $\sigma=0.4$, the actual reward for Node A alone is:

$$58,084 \times 0.4 + 0.6 \times \frac{3000}{4000} = 49371$$

As Node B does not have any delegator, so it will retain 100% of its rewards (i.e. 50,539).

As we assume δ is 0.1, Reward A is 10% of the daily rewards for training nodes, while 90% of the rewards (i.e. Reward B) will be locked until the end of task. Thus, Reward A and Reward B for Node A on Day 1 are 4,937.1 and 44,433.9 respectively, whereas those for Node B are 5,053.9 and 45,485.1 respectively.

As for validators, for the same task on Day 1, total rewards for all validators are:

$$309157.68 \times \frac{12000}{6500 + 12000} = 200534.0$$

Then, if we assume Validator A has a score of 0.472768, the the reward for validator A only is:

$$200534.0 \times 0.369 \times 0.4 + 0.6 \times \frac{3000}{3000+0} = 73997$$

Similarly, on Day 1, rewards for Validator B (score = 0.280226, stake = 6,000) and Validator C (score = 0.247006, stake = 3,000) are 87,834 and 38,703 respectively.

As for Delegator A who delegates 1,000 to Node A, its rewards on Day 1 are calculated as such:

$$58084 \times 0.6 \times \frac{1000}{3000+1000} = 8712.6$$

The reward distribution among the participants in task 1 during the 5 days is shown in Table II.

VI. ODIN AI CONSENSUS IN FL ALLIANCE

Figure 5 depicts the workflow of a FL Alliance task in ODIN AI. As shown in our work in leveraging blockchain to defend against poisoning attacks in FL Alliance [14], ODIN AI adopts a distributed voting and a reward-and-slash mechanism to construct secure FL Alliance systems.

A. Task Creators

Similar to task creation in AI Arena, an FL Alliance task creator must satisfy predefined criteria. Only FL Alliance tasks verified by the ODIN AI DAO will be eligible for rewards from the daily emissions. Otherwise, the FL Alliance task creator must self-fund the reward pool using their own $\$ODIN AI$.

B. Random Role Selection

Consider a FL Alliance task involving P participants, denoted as $P = 1, \dots, P$. To participate in the training process, each participant needs to stake a specified quantity of coins. Upon formally joining the training task, each participant's local dataset D_p is randomly partitioned into a training set D_p^{train} and a test set D_p^{test} , which will not be shared with other participants at any time. At the beginning of each round t in the FL Alliance task, participants are randomly assigned roles as either a proposer (P_T) or a voter (P_V) through an on-chain random function. Subsequently, a model initialised or pre-trained model downloaded from AI Arena by one of the proposers is chosen at random temporarily to serve as the pioneering global model. The selected model's weights or gradients are then distributed to all participants, ensuring a unified starting point for local models. Proposers are responsible for training their local models using their own data and subsequently sharing the updated model weights

Algorithm 1 ODIN AI Federated Training.

T: Total number of global communication rounds
 E: Total number of local model update epochs
 θ_g : Global model
 D_p^{train} : local training dataset
 D_p^{test} : local test dataset

- 1: **procedure** INIT
- 2: Download pre-trained model from AI Arena or initialise global model θ_g^t
- 3: Broadcast θ_g^t to all participants
- 4: **end procedure**
- 5: **procedure** UPDATE(θ_g^t)
- 6: Initialise local model θ_p^t
- 7: Update model parameters $\theta_p^t \leftarrow \theta_g^t$
- 8: $\theta_p^{t+1} \leftarrow \theta_p^t - \eta \nabla L(\theta_p^t; b)$ ▷ Local model update
- 9: **return** θ_p^{t+1}
- 10: **end procedure**
- 11: **procedure** EVALUATE($\theta_t + 1^p$)
- 12: $\theta_g^{t+1} \leftarrow \sum_{p=1}^P \frac{n_p}{N} \theta_p^{t+1}$ ▷ Model updates aggregation
- 13: Res $\leftarrow \frac{\theta_g^{t+1}}{D_p^{test}}$ ▷ Evaluate aggregated model
- 14: **return** Res
- 15: **end procedure**
- 16: **procedure** MAIN
- 17: Random select initialisation leader $\rightarrow p_{lead}$
- 18: $\rightarrow p_{lead}$ Do procedure **Init**
- 19: **for** t = 1 **to** T **do**
- 20: Random assign roles for all participants
- 21: Proposer does procedure **Update**
- 22: Voter does procedure **Evaluate**
- 23: $vote_p^t \leftarrow \text{calculate } Res_p^t$
- 24: $aggVote^t = \sum_{p=1}^{P_v} vote_p^t$ ▷ Votes aggregation
- 25: $\theta_g^{t+1} \leftarrow \text{GlobalModelSelection}(aggVote^t)$
- 26: Broadcast θ_g^{t+1}
- 27: **end for**
- 28: **end procedure**

or gradients with all participants. Voters, on the other hand, aggregate these updates from proposers. They then proceed to validate the aggregated model updates, resulting in the generation of a validation score.

C. FL Alliance Training

At the start of round t, proposers initially download the global model, denoted as θ^{t-1} , which was finalised in the previous round (t-1). Using the local training dataset D_p^{train} , proposers then proceed to update the model θ^{t-1} through E epochs of local training. Then the updated model θ^t of the current round t will be uploaded to the voters for evaluation.

D. FL Alliance Aggregation

Upon the completion of task training during round t, the voter gathers the local models $\{\theta^t\}_{p=1}^{P_T}$ from proposers. These

models are then aggregated into the latest global model using a weighted averaging approach, as described below:

$$\hat{\theta}^t = \sum_{p=1}^{P_T} \beta_p \cdot \theta_p^t$$

Here, the weight β_p is defined as $\frac{n_p}{N}$, with $n_p = |D_p^{train}|$ indicates the number of local training data samples for each proposer p, and $N = \sum_{p=1}^{P_T} n_p$ is the total number of training data samples across all proposers P_T .

E. FL Alliance Validation and Voting

- After the model aggregation process is finalised, the voter proceeds to evaluate the aggregated model $\hat{\theta}^t$, utilising their own local testing datasets D^{test} . This evaluation phase involves the computation of a local validation

Algorithm 2 Reward-and-slash design for FL clients.

P: Set of participants at round t
rewardPool: Total reward pool
Bal_k: Stake amount of participant k
s_p: Slashed rate
participantsDistributionRate: Ratio of participants in one round
TotalRounds: Total numbers of communication round
BalThreshold_p: minimum stake of participants

```

1: poolp ← initialRewardPoolSize
2: for  $t = 0; t < \text{TotalRounds}$  do
3:   Candidates ← {}
4:   for  $k \in P$  do
5:     if  $\text{Bal}_k \geq \text{BalThreshold}_p$  then
6:       Candidates ← Candidates  $\cup \{k\}$ 
7:     end if
8:   end for
9:   Ppt, Pvt ← Randomly select proposers and voters from Candidates based on participantsDistributionRate
10:  aggVotet ← Aggregated voting results from the voters Pvt at round  $t$ 
11:  roundRewardAmount ←  $\frac{\text{pool}_p}{\text{TotalRounds} - t}$ 
12:  roundTotalStakedTokensForGoodParticipants ← 0
13:  if  $\text{aggVote}^t \geq 0$  then ▷ Compute the participants that should be rewarded in this round
14:    for  $k \in P_p^t$  do
15:      roundTotalStakedTokensForGoodParticipants += Balk
16:    end for
17:  end if
18:  for  $k \in P_v^t$  do
19:    if  $\text{vote}_k^t \cdot \text{aggVote}^t \geq 0$  then
20:      roundTotalStakedTokensForGoodParticipants += Balk
21:    end if
22:  end for
23:  if  $\text{aggVote}^t \geq 0$  then ▷ Reward and Slash Proposers
24:    for  $k \in P_p^t$  do
25:      Balk ← Balk + roundRewardAmount ·  $\frac{\text{Bal}_k}{\text{roundTotalStakedTokensForGoodParticipants}}$ 
26:      poolp ← poolp - roundRewardAmount ·  $\frac{\text{Bal}_k}{\text{roundTotalStakedTokensForGoodParticipants}}$ 
27:    end for
28:  else
29:    for  $k \in P_p^t$  do
30:      Balk ← Balk - Balk · sp
31:      poolp ← poolp + Balk · sp
32:    end for
33:  end if
34:  for  $k \in P_v^t$  do ▷ Reward and Slash Voters
35:    if  $\text{vote}_k^t \cdot \text{aggVote}^t \geq 0$  then
36:      Balk ← Balk + roundRewardAmount ·  $\frac{\text{Bal}_k}{\text{roundTotalStakedTokensForGoodParticipants}}$ 
37:      poolp ← poolp - roundRewardAmount ·  $\frac{\text{Bal}_k}{\text{roundTotalStakedTokensForGoodParticipants}}$ 
38:    else
39:      Balk ← Balk - Balk · sp
40:      poolp ← poolp + Balk · sp
41:    end if
42:  end for
43:  t ← t + 1
44: end for

```

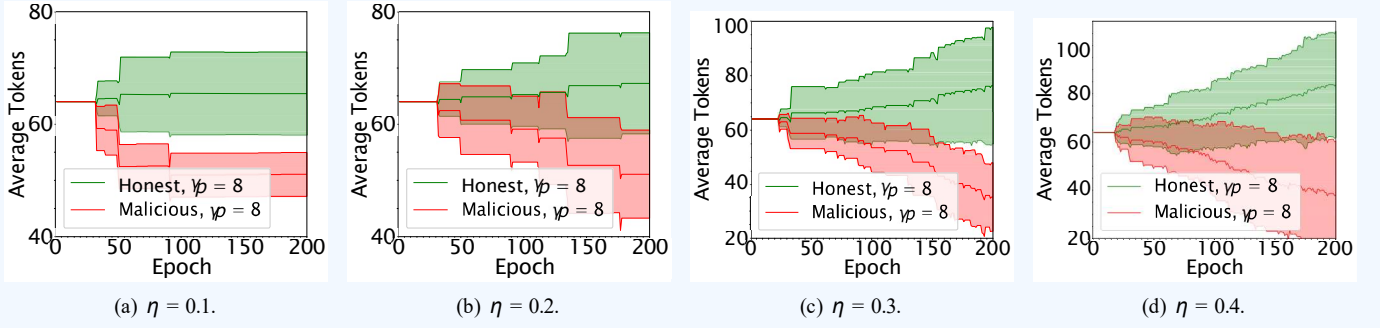


Fig. 6: Example: FL system with reward and slash mechanism under different values of the ratio of malicious clients (η), taken from [14]. The average balance of honest clients increases, while the average balance of malicious clients decreases over time.

score, s_p^t , which functions as a criterion for assessing the model's performance. These individual validation scores are then submitted to a smart contract for aggregation. Following the aggregation, the aggregated score is compared with the previous round's score, s_p^{t-1} , to assess progress or decline in model performance. The smart contract then determines the next steps for the aggregated model based on these scores: advancement to the next phase for satisfactory performance improvement, or a return to the preceding validated model to begin a new cycle of training, aggregation and evaluation, if progress is deemed insufficient.

$$\text{vote}_p^t = \begin{cases} 1, & s_p^t \geq (1 - \epsilon) \cdot s_p^{t-1} \\ -1, & s_p^t < (1 - \epsilon) \cdot s_p^{t-1} \end{cases}$$

Here, ϵ is a hyperparameter within the range (0, 1), designated to tolerate the permissible margin of performance decline across successive rounds.

After receiving all reported voting results $\{\text{vote}_1^t, \dots, \text{vote}_{P_v}^t\}$ from the validators, the aggregator will calculate the aggregated voting result via the following formula:

$$\text{aggVote}^t = \sum_{p=1}^{P_v} \text{vote}_p^t$$

For each round t , the finalised aggregated global model update is determined by the aggregated voting result:

$$\theta^t = \begin{cases} \hat{\theta}^t, & \text{aggVote}^t \geq 0 \\ \theta^{t-1}, & \text{aggVote}^t < 0 \end{cases}$$

F. FL Alliance Rewards for Participants

The aggregated voting result aggVote^t will also determine the rewards distribution for participants in a FL Alliance task.

- **Rewards and Penalties for Proposers/Training Nodes:** As shown in Algorithm 2, in any given round t , should aggVote^t be non-negative, all training nodes selected for that round will receive rewards. Conversely, a negative aggregate vote will result in penalties for these nodes.

- **Rewards and Penalties for Voters/Validators:** As shown in Algorithm 2, for round t , if $\text{aggVote}^t \geq 0$, then validators who issued a positive vote will be rewarded, while others will face penalties. Conversely, should the aggregate vote be negative, validators who aligned with this outcome are rewarded, whereas those who did not will be penalised.

G. Example

As illustrated in Figure 6(d), taken from our previous work [14], proper configuration of the slashing and reward mechanisms enables the expulsion of malicious FL participants from the system, while incentivising honest behaviour.

H. FL Alliance Improvement: ZKPs-based FL

ODIN AI also adopts advanced techniques such as Zero-knowledge proof (ZKP) [16], [17] to construct secure decentralised AI training systems.

ZKPs for FL Alliance Aggregation: As demonstrated in our prior study, ODIN AI incorporates ZKP to address

issues arising from the centralisation of the FL Alliance aggregator/server, as detailed in our earlier research [18]. Our FL system, which can be underpinned by both blockchain technology and ZKPs and function in the following manner:

- **Setup Phase:** Each participant, comprising N clients and an aggregator, generates their unique private/public key pairs. These pairs are directly associated with their respective blockchain addresses.
- **Client Selection Phase:** At the beginning of each epoch, a subset of n clients is selected from the total N by using Verifiable Random Functions.
- **Local Computation Phase:** The selected n clients start local model training to derive their individual model updates w_1, w_2, \dots, w_n . Utilising the Pedersen commitment, each client encrypts their update as $\text{Enc}(w_i) = g^{w_i} \cdot h^s$, where g and h are predefined public parameters and s_i is a randomly generated number by the client. Following encryption, clients authenticate these updates using their private keys to produce a signature sig and subsequently transmit the compilation of their local

model update, the generated random number, the encrypted update, and the signature $(w_i, s_i, \text{Enc}(w_i), \text{sig}_i)$ to the aggregator.

- **Aggregation and ZKP Generation Phase:** The aggregator aggregates the incoming local updates to form a unified global model update $w = \sum_{i=1}^n w_i$. It also calculates the collective encrypted value of this global update as $\text{Enc}(w) = \prod_{i=1}^n \text{Enc}(w_i)$ and signs this encrypted value to produce a signature sig . Utilising zkSnarks, the aggregator issues a proof π to validate the accuracy and authenticity of the aggregation process, based on the provided statement and witness, ensuring the integrity of both the individual updates and the aggregate model. Specifically, the aggregator then leverages zkSnark to issue a proof π for the following statement and witness:

$$\begin{aligned} \text{statement} &= (\text{Enc}(w_1), \text{sig}_1, \text{Enc}(w_2), \text{sig}_2, \\ &\quad \dots, \text{Enc}(w_n), \text{sig}_n, \text{Enc}(w)) \\ \text{witness} &= (w_1, s_1, w_2, s_2, \dots, w_n, s_n, w) \end{aligned}$$

where the corresponding circuit $C(\text{statement}, \text{witness})$ outputs 0 if and only if:

$$\begin{aligned} \forall 1 \leq i \leq n, \text{Enc}(w_i) &= g^{w_i} \cdot h^{s_i} \\ w &= \sum_{i=1}^n w_i \\ \text{sig}_i &\text{ is signed by the client } i \end{aligned}$$

- **Global Model and Proof Dissemination Phase:** The aggregator distributes the global model update w and its encryption $\text{Enc}(w)$ back to the n clients. Concurrently, it broadcasts the validity proof π along with the encrypted global model update to the block proposers.
- **Blockchain Verification Phase:** Upon receiving the proof π and the encrypted global model update from the aggregator, block proposers verify π . If deemed valid, the hash of $H(\text{Enc}(w))$ is inscribed onto the blockchain, cementing the update's correctness.
- **Blockchain Consultation Phase:** As a new epoch initiates, the next cohort of n selected clients peruses the blockchain to verify the inclusion of $H(\text{Enc}(w))$. Upon successful validation, they proceed with their local training, guided by the insights gleaned from the aggregated global model update w .

VII. ODIN AI GOVERNANCE

ODIN AI token holders are entitled to engage in the system's democratised governance through a DAO. To participate in governance, token holders typically need to lock their tokens in a smart contract. Each token can represent a vote, aligning the distribution of power proportional to users' stake.

Users can propose, debate, and vote on various aspects of development and management, from technical updates and protocol modifications to treasury management and community initiatives.

- **Proposing:** The ODIN AI community actively shapes the protocol's future through a proposal system for all token

holders. Proposals can range from addressing technical issues like bug fixes and algorithm optimisation to driving wider community impact, such as allocating treasury funds for research or launching educational programs.

- **Debating:** Proposed ideas are then open for discussion and critique within the ODIN AI community. Token holders can engage in forums, discussions, and possibly even direct communication with developers to analyze the merits and potential consequences of each proposal. This debate fosters transparency and ensures that decisions are well-informed and considered from multiple perspectives.
- **Voting:** Once a proposal has been sufficiently debated, token holders cast their votes to decide its fate. The voting system likely incorporates mechanisms like weighted voting (where larger holdings carry more weight) or quadratic voting (which incentivises thoughtful contributions and discourages manipulation) to ensure fair representation.

The statement emphasises that ODIN AI's governance model allows for continuous adaptation as the platform and the decentralised AI landscape evolve:

- **Policy Adaptation:** As new challenges and opportunities arise, token holders can use the voting system to modify existing policies or create entirely new ones. This ensures that ODIN AI remains relevant and responsive to the changing needs of its community and the broader AI ecosystem.
- **Feature Implementation:** Proposals for implementing new features can be put forward and voted on, allowing the ODIN AI platform to grow and evolve based on user demand and feedback. This fosters innovation and keeps ODIN AI at the forefront of decentralised AI development.
- **Responding to Challenges:** The ability to quickly adapt policies and implement changes allows ODIN AI to effectively respond to unforeseen challenges like security vulnerabilities, regulatory shifts, or market fluctuations.

As ODIN AI and decentralised AI landscape mature, token holders can adapt policies, implement new features, and respond to emerging challenges.

VIII. ODIN AI APPLICATIONS

The ODIN AI system can be used to construct centralised AI, which have been proven to applied in the following cases.

A. Decentralised AI for LLMs

- **Pre-training of LLMs:** ODIN AI facilitates the pre-training of LLMs by leveraging a decentralised network whereby members can contribute computational resources and diverse data sets. This unlocks proprietary data that would otherwise remain inaccessible or unused in traditional, centralised open-source development. Diverse datasets ensure LLM versatility and ensures a broader representation of linguistic and cultural nuances, as well as community-defined values for LLMs.
- **Fine-tuning of LLMs:** Fine-tuning involves adapting a pre-trained model to perform specific tasks or improve

its accuracy on particular types of data. ODIN AI supports fine-tuning in several ways:

- *Fine-tuning for Financial Transactions:* LLMs can be fine-tuned to act as intelligent agents for cryptocurrency transactions. Capabilities include transfers, swaps, and bridging between different cryptocurrencies. ODIN AI’s planned collaborations with platforms such as [Morpheus Network](#) and [Oxscope](#) can facilitate hosting these AI models, ensuring that they are accessible and operational for the community. This enables secure and efficient AI-driven financial transactions.
- *Fine-tuning for AI Companions:* AI models can be fine-tuned to interact with users in more personalised and engaging ways, similar to those on platforms like [Character.ai](#). ODIN AI can host these sophisticated AI companions, enhancing user experience through more natural and context-aware interactions.

B. Decentralised AI for Stable Diffusion Models

The ODIN AI system can be used to fine-tune Stable Diffusion text-to-image models. One critical component of this process involves Low-Rank Adaptation (LoRA) [19], which modifies certain parameters within the model’s architecture to make it more adaptable to specific tasks without extensive retraining.

- **Fine-tuning LoRA:** LoRA is designed to adapt pre-trained models by introducing trainable low-rank matrices into the architecture. This technique allows for efficient adaptation with minimal additional computational cost and a smaller number of trainable parameters. In the context of ODIN AI and Stable Diffusion Models, applying LoRA is particularly advantageous for several reasons:
 - *Community-Driven Enhancements:* By decentralising the fine-tuning process, ODIN AI broadens participation in contributing specific knowledge and preferences. Artists, designers, and other creatives can input unique styles or features they wish to see enhanced, improving output quality and ensuring that it serves a wider array of cultural contexts and artistic expressions.
 - *Scalability and Accessibility:* Fine-tuning with LoRA can be scaled across multiple nodes, facilitating more widespread and continuously iterative improvements.
 - *Use Case Expansion:* By fine-tuning Stable Diffusion Models with LoRA, ODIN AI can cater to specific industries or niches. For example, the model could be fine-tuned to generate medical illustrations for educational purposes, architectural visualisations for real estate, or unique art styles for digital media.

C. Decentralised AI for Linear Regression Models

Linear regression models [20] are fundamental tools in statistical analysis and predictive modeling, widely used for their simplicity and effectiveness in understanding relationships between variables. ODIN AI applies these principles in a decentralised setting to address specific healthcare challenges, such as diabetes management.

Diabetes management presents a critical area where linear regression can be effectively utilised to predict patient outcomes based on various inputs such as blood sugar levels, diet, exercise, and medication adherence. FL Alliance facilitates the development of these predictive models with decentralised data sources in a way that respects patient data protection.

- *Data Protection and Security:* ODIN AI allows multiple healthcare providers to collaborate in the model training process without actually sharing the data. This method is crucial for complying with stringent health data protection regulations such as HIPAA in the U.S. Each participant (e.g., hospitals, and clinics) retains control over their data, which is used to compute model updates locally. These updates are then aggregated to improve a shared model without exposing individual patient data.
- *Enhanced Model Accuracy and Reliability:* By integrating data from a diverse range of demographics and geographical locations, ODIN AI can help develop more accurate and generalised linear regression models for diabetes management. This diversity is especially important in healthcare, where patient populations can vary significantly, affecting the reliability of predictive models.
- *Collaborative Innovation:* Different healthcare entities contribute to a common goal, accelerating innovation and leading to the discovery of novel insights into diabetes management and treatment strategies.

IX. CONCLUSION

ODIN AI provides solutions to build decentralised AI through AI Arena, FL Alliance, and AI Marketplace. ODIN AI dismantles obstacles that hinder participation in AI systems, enabling developers to contribute models, data, or computational resources in a flexible, modular fashion. ODIN AI fosters the creation of a diverse array of models, meticulously crafted by and expressly for the communities they serve in AI models.

REFERENCES

- [1] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008.
- [2] G. Wood, “Ethereum: A secure decentralised generalised transaction ledger,” *Ethereum project yellow paper*, vol. 151, pp. 1–32, 2014.
- [3] S. Werner, D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. Knottenbelt, “Sok: Decentralized finance (defi),” in *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 2022.
- [4] P. Bellavista, L. Foschini, and A. Mora, “Decentralised learning in federated deployment environments: A system-level survey,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 1, pp. 1–38, 2021.
- [5] T. Wu, S. He, J. Liu, S. Sun, K. Liu, Q.-L. Han, and Y. Tang, “A brief overview of chatgpt: The history, status quo and potential future development,” *IEEE/CAA Journal of Automatica Sinica*, 2023.
- [6] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang, “Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [7] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen *et al.*, “Palm 2 technical report,” *arXiv preprint arXiv:2305.10403*, 2023.
- [8] “Why google’s ai tool was slammed for showing images of people of colour,” available at: <https://www.aljazeera.com/news/2024/3/9/why-google-gemini-wont-show-you-white-people>.
- [9] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong *et al.*, “A survey of large language models,” *arXiv preprint arXiv:2303.18223*, 2023.

- [10] NIST, “There’s more to ai bias than biased data, nist report highlights,” available at: <https://www.nist.gov/news-events/news/2022/03/theres-more-ai-bias-biased-data-nist-report-highlights>.
- [11] IBM, “Shedding light on ai bias with real world examples,” available at: <https://www.ibm.com/blog/shedding-light-on-ai-bias-with-real-world-examples/>.
- [12] “When ai gets it wrong: Addressing ai hallucinations and bias,” available at: <https://mitsloanedtech.mit.edu/ai/basics/addressing-ai-hallucinations-and-bias/>.
- [13] N. Dong, J. Sun, Z. Wang, S. Zhang, and S. Zheng, “ODIN AI : Defending malicious behaviors in federated learning with blockchain,” *NeurIPS 2022 Workshops on Decentralization and Trustworthy Machine Learning in Web3: Methodologies, Platforms, and Applications. Runner-up Award.*, 2022.
- [14] N. Dong, Z. Wang, J. Sun, M. Kampffmeyer, W. Knottenbelt, and E. Xing, “Defending against poisoning attacks in federated learning with blockchain,” *IEEE Transactions on Artificial Intelligence*, 2024, <https://arxiv.org/pdf/2307.00543.pdf>.
- [15] 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*. PMLR, 2017.
- [16] E. B. Sasse, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, “Zerocash: Decentralized anonymous payments from bitcoin,” in *IEEE Symposium on Security and Privacy*, 2014.
- [17] J. Groth, “On the size of pairing-based non-interactive arguments,” in *Advances in Cryptology—EUROCRYPT*, 2016.
- [18] Z. Wang, N. Dong, J. Sun, W. Knottenbelt, and Y. Guo, “zkfl: Zero-knowledge proof-based gradient aggregation for federated learning,” *IEEE Transactions on Big Data*, 2024.
- [19] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “Lora: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [20] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to linear regression analysis*. John Wiley & Sons, 2021.