International Journal for Innovative Engineering and Management Research
PEER REVIEWED OPEN ACCESS INTERNATIONAL JOURNAL
www.ijiemr.org

AI Safety, Human Control and Alignment: Decentralised AI System Using Holochain

Ms. Afifa Mubeen
Student, BS MS Computer Science, Vishwa Vishwani Institute Of Systems and Management.
Hyderabad, Telangana, India | afifamubeen26@gmail.com | 9652304795
Ms. Madeeha Mubeen
Student, BS MS Computer Science, Vishwa Vishwani Institute Of Systems and Management.
Hyderabad, Telangana, India | madeehamubeen14@gmail.com | 8074822721
Ms. Mariah Ashfaq
Student, CSE, Shadan Women's College of Engineering and Technology.
Hyderabad, Telangana, India | mariahashfaq06@gmail.com | 7093690082
Ms. Nahdiya Fatima Younus
Student, CSE, Shadan Women's College of Engineering and Technology.
Hyderabad, Telangana, India | nahdiyayounus@gmail.com | 6302752895

ABSTRACT:

Artificial Intelligence has been predominantly built on centralized systems, with all power of computing, managing and storage relying upon indefinite single institutions. In its entirety centralized systems propose conflict in Privacy and Data Security, Loss of Data Sovereignty and Monopolistic agency. To counter these we introduce Decentralised AI systems which are distributed, autonomous systems that work on individual entities and do not have a central computational unit. Instead, the AI governance is outsourced to multiple entities that are binded in a network reducing its intensive abuse of power and control due to distributed data flow and computing. Such a workflow is prevalent in preventing data misuse. The system proposed here uses a holochain technology that allows peer to peer networking for light weight communication along with extensions like Akash and Render for heavy computing. Holochain is a budding technology that promises Efficiency, Scalability and Data Sovereignty. It is primarily agent centric, i.e. multiple independent source chains per agent along with shared a DHT (Distributed Hash Table). Each node only stores its own data rather than a full ledger, and a shard of the DHT. In decentralized artificial intelligence systems, responsibility is spread out and reviewers review different parts with coordinated escalation. While the authority could be shared between developers and governing bodies.

Decentralized AI is an alternative to centralized AI and holochain technology facilitates in enticing it, making it more efficient, scalable and ensuring data security. Distributed control enables more localized agency and system autonomy, evidently making the system strong, reliable and secure.

Keywords: Decentralized AI, Holochain, Data Sovereignty, Distributed Computing, Agent Centric Architecture, Data Privacy and Security

1. INTRODUCTION:

1.1 Centralized AI

Artificial Intelligence (AI) has been explosive in many fields and subfields that differ from academics to workspaces. Most, if not all, rely on a centralized computing system that is responsible for managing, processing and governance. (Vipra & West, 2023; V. Kersic and M. Turkanovic., 2026). Centralized AI raises questions on data privacy and integrity due data pooling for training and deployment on central servers, this concerns data breaches and misuse of available data. (Carlini et al., 2021; Shokri et al., 2017). Heavy dependency on individual providers can cause issues related to vendor lock in, could limit transparency and reduce user control over data and flow of data (Opara-Martins et al., 2016; Whittaker, 2021). This paradigm concludes centralized AI to be a rather risky structure to govern and maintain ethical boundaries.

1.2 Decentralized AI:

A possible alternative is decentralized AI, which spreads data storage, training and validation across independent entities instead of a single entity.(Montes & Goertzel, 2019). Methods such as federated learning keep the raw data in the local devices and only upload the model changes (McMahan et al., 2017). Moreover, the combination of AI with distributed ledgers has been proposed to enhance trust, traceability and incentives (Dinh & Thai, 2018; Salah et al., 2019). This direction is of interest to researchers for its promise of stronger privacy, greater resilience and more transparent governance although the practical architecture remains an open question.

1.3 Proposed Technology:

The framework of Holochain is agent centric, where each agent keeps their own signed source chain and exchanges data using a validated distributed hash table, rather than requiring all nodes to agree on a single global ledger (Harris-Braun et al., 2018). We chose Holochain in favor of traditional blockchain since consensus across the globe has serious consequences in terms of scalability, latency, and energy use (de Vries, 2018; Zheng et al., 2018), while with

Holochain, applications can scale up according to the number of its users. However, Holochain apps run on end users' own gadgets and are not designed to have enough computing power available in form of GPUs required for training and running modern AI algorithms (Holochain, n.d.). This is the gap that leads us to Akash (Akash Network 2018), a decentralised marketplace for cloud computing and Render (Render Network 2020), a distributed GPU rendering and compute network. Both networks can offer on demand GPU power, without creating another centralized service again.

The present paper describes a conceptual framework where Holochain, Akash, and Render come together in order to create a decentralised artificial intelligence. The current work highlights benefits of this integration, as well as the corresponding limitations.While technical constraints include coordination cost, verifying outsourced computations, and latency issues. At the same time, the governance aspects of the framework, where multiple agents take different actions independently, have been covered. The paper does not propose a working framework but only the one for future research.

2. LITERATURE REVIEW:

2.1 Centralised and Decentralised AI Systems:

Modern AI relies on a small number of firms that dominate compute power, data and research agendas (Vipra & West, 2023; Whittaker, 2021) and unmonitored data flow can result in control being in the hands of big technology companies (Widder et al., 2023). The centralization also creates points of failure, with pooled data and models open to extraction of training data from language models (Carlini et al., 2021), and copying of models via their prediction interfaces (Tramèr et al., 2016). Decentralized AI distributes control (Montes & Goertzel, 2019); federated learning keeps raw data on local devices and only shares model updates (McMahan et al., 2017), though privacy leakage from the model updates is a concern (Li et al., 2020).

2.2 Blockchain and Holochain:

To facilitate trust, traceability and incentives, AI researchers have combined the two paradigms of distributed ledgers (Dinh & Thai, 2018; Salah et al., 2019). Examples of such systems include DeepChain, a blockchain based incentive and audit mechanism for privacy preserving

collaborative training (Weng et al., 2021) and a federated learning scheme, which tracks exchanges of industrial data on a blockchain (Lu et al., 2020). However, global consensus reduces throughput and introduces latency and proof of work based chains are energy intensive (de Vries, 2018; Nakamoto, 2008), whereas Wüst and Gervais (2018) argue that most blockchain applications have no need for one. Holochain, takes a different approach, whereby each agent stores a signed, append only source chain, publishes certain data onto a Distributed Hash table and depends on others to verify it against common rules (Harris-Braun et al., 2018).

2.3 Decentralised Computing: Akash and Render

Training and serving modern AI typically requires large amounts of GPU compute, concentrated among a few firms (Vipra & West, 2023), and neither blockchains, nor holochains provide it. Decentralised compute marketplaces attempt to fill the void. Akash is an open marketplace where providers bid to host containerised workloads and tenants select from the bids (Akash Network, 2018). Render started as a network for distributed GPU rendering, and has expanded into AI inferencing through the Render Compute Network (Render Network, 2020). Akash hosts generic containerised workloads, whereas Render focuses on GPU intensive jobs. Also, outsourcing invites the question of correctness of results; that is, verifiable computation allows a client to verify a result without having to re execute it, but practical overhead remains a challenge (Walfish & Blumberg, 2015). Evidence about both networks comes mostly from the projects' documents rather than independent studies.

3. Research Gap:

The reviewed literature demonstrates that each level of the decentralised AI was developed in isolation from the others. Federated learning keeps data local, blockchain schemes adds trust and incentives, Holochain proposes an agent centric coordination without global consensus, and Akash and Render offer decentralised GPU compute. What is left unsolved is the question of composition: how agent centric coordination would interact with decentralised cloud and GPU infrastructure, how responsibilities would be divided between Holochain, Akash and Render, and how information and computational results would be exchanged between them while being verifiable at a practical cost.

Apart from the engineering challenges, allocating tasks between independent providers requires addressing some governance issues. Outsourcing AI tasks to unknown providers involves risks to privacy, access to results, and responsibility for defective or intentionally harmful outcomes and the importance of human agency. Studies on Akash and Render that have been made until now mostly concern technical and economic aspects of the technology and rely mostly on information taken from the companies' documents. Consequently, it may be concluded that, in addition to the technical feasibility study, the issue of the governance model in such systems should also be considered.

4. Objectives of the Study:

Aim: To examine, at a conceptual level, whether a decentralised AI system based on Holochain, with Akash and Render providing GPU compute, addresses the privacy, data sovereignty, concentration and safety challenges of centralised AI and what constraints it would encounter.

Specific objectives:

1. To review evidence on the privacy, security and governance challenges of centralised AI.

2. To describe how Holochain, Akash and Render each function and what capabilities they do and do not offer.

3. To propose a layered conceptual architecture in which Holochain is responsible for coordination and data ownership and Akash and Render provide heavy computation.

4. To analyse the technical constraints of the design, particularly verification of outsourced computation, and coordination cost.

5. To analyse the implications for safety, human control and alignment of such a distributed architecture and recommend governance measures and directions of validation.

5. Research Methodology:

5.1 Research Design:

The nature of our study is qualitative and conceptual. In our study, a conceptual architecture is proposed as a solution to the identified problem and subsequently evaluated in relation to the specified criteria. Since this paper does not involve construction or implementation of a working system, the contribution is more in terms of reasoned design and analytical evaluation than in empirical findings, which fits within the scope defined in the Introduction.

5.2 Sources and selection:

Evidence comes from three types of source: peer reviewed papers and conference articles on privacy in AI, federated learning, blockchain based AI, verifiable computation and AI safety; preprint articles and white papers for technologies which are not yet subject to peer review, notably Holochain; and official documentation and project reports for Akash and Render. The review is organised by concept, as recommended by Webster and Watson (2002) Any claim about Akash or Render depends on material published by the project itself; this has been stated since such material has not been subject to independent peer review.

5.3 Analytical procedure:

The study follows the stages indicated in Figure 1: stating the problem, conducting a literature review and identifying a gap, identifying objectives, individually analysing the three technologies, designing a combined architecture, and evaluating that architecture against the criteria in Section 5.4. The findings and recommendations thus emerge, and their recommendations point towards the verification that this conceptual paper cannot itself

provide.

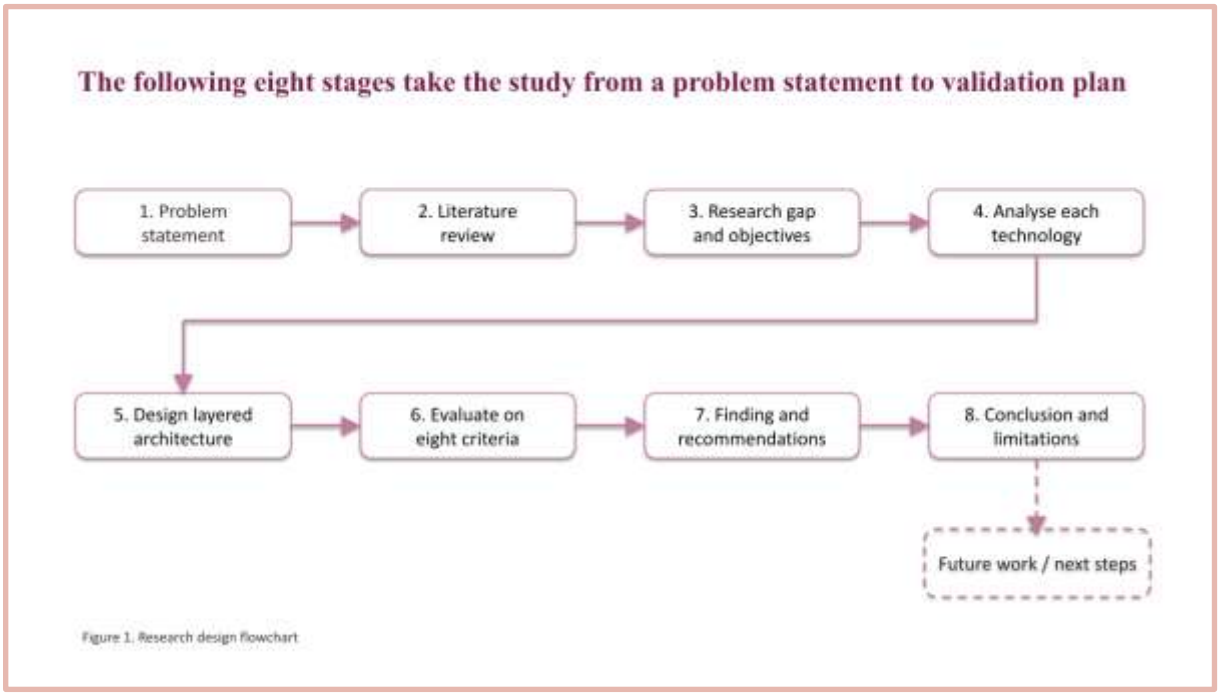


Figure 1. Research design flowchart

5.4 Evaluation Criteria: The architecture is compared on the basis of the following criteria.

Criterion	What it asks
Data privacy	Can the raw data stay with its owner?
Data sovereignty	Does the owner control use and deletion?
Scalability	Does capacity scale with the number of participants?
Compute capability	Can it run both training and inference on GPUs?
Verifiability	Can the results be verified?
Resource efficiency	What is the energy and cost overhead?
Governance and oversight	Can humans set the rules, audit them and change them?

5.5 Validity and limitations:

The biggest limitations are that there is no prototype. Evidence on Holochain, Akash and Render was largely self published, and technologies are changing rapidly, making details

potentially out of date. No human participants or personal data were used, and therefore no ethical approval was needed.

6. Conceptual Background:

6.1 Holochain: agent centric coordination:

Holochain is based on an agent based architecture and does not use a centralized global ledger. Every agent stores their actions in a cryptographically signed, append only source chain. Chosen transactions are then committed to a Distributed Hash Table (DHT), where responsible nodes store and validate them based on the set of application defined rules (Harris-Braun et al., 2018). Unlike traditional blockchains, Holochain does not use consensus across the entire network and uses distributed peer validation instead. Malicious transactions can be recognized through validation methods with signed warrants serving as proof of rule breaking (Holochain, n.d.-b). As per the discussed architecture, Holochain would enable decentralized coordination through storing consent records, access rights, model identifiers and computational task information. Nevertheless, Holochain is unable to offer dedicated GPU hardware necessary for computationally heavy AI tasks (Holochain, n.d.-a).

6.2 Akash Network: decentralised cloud computing:

Akash is a decentralized platform for cloud computing. The tenant sets up the resource requirements and a container image in the configuration file, while providers bid through a reverse auction, and the tenant chooses a bid, setting up a lease. The provider executes the workload and gets payment out of the tenant's deposit (Akash Network, n.d.). The containerized nature of workloads makes it appropriate for long term services and generic AI tasks such as model serving, and it can provide GPU capabilities. This platform can be used for hosting AI services such as model serving endpoints. Its major drawback is that it distributes trust across independent providers but does not ensure correct execution or privacy of data.

6.3 Render Network: distributed GPU processing:

Render matches creators who require their GPU needs to be met by operators who provide their idle GPUs, rewarding them with tokens from the network (Render Network, 2020). Originally

intended for 3D rendering tasks, it has been broadened to accommodate AI tasks using a compute network that targets inference and edge machine learning (Render Network Foundation, n.d.). Akash focuses on general hosting services; Render, on the other hand, provides huge amounts of consumer GPUs for batch and inference tasks, and acts as the second compute resource.

6.4 Life Cycle of a Decentralized AI system:

A decentralised AI system consists of four stages that can be distributed among different agents: data storage, training, validation and inference (Montes & Goertzel, 2019). The proposed division of work is as follows:

Stage	Description	Location
Data ownership and consent	Data owners retain raw data and provide permission and metadata information	Holochain: source chain of the owner and DHT
Training	Local/Federated Training; Outsourcing of intense rounds to be done by GPUs	Owners' devices for local training; Akash/Render for heavy processing
Validation	Peer nodes verify updates, results, and attestations against common standards	Holochain validation rules
Inference	Model delivery to users; Sensitive questions are answered locally where possible	Akash/Render for model delivery; Holochain for access control and audit

6.5 Proposed Layered Architecture:

International Journal for Innovative Engineering and Management Research
PEER REVIEWED OPEN ACCESS INTERNATIONAL JOURNAL
www.ijiemr.org

The figure 2 below puts together the components in four layers. The participants and data remain at the edge, the Holochain layer coordinates, verifies, and records; Akash and Render layer provide heavy computations on demand, while a governance layer is there to lay out the rules and deal with escalations. This design is purely conceptual for the purpose of analysis and does not claim an implemented system.

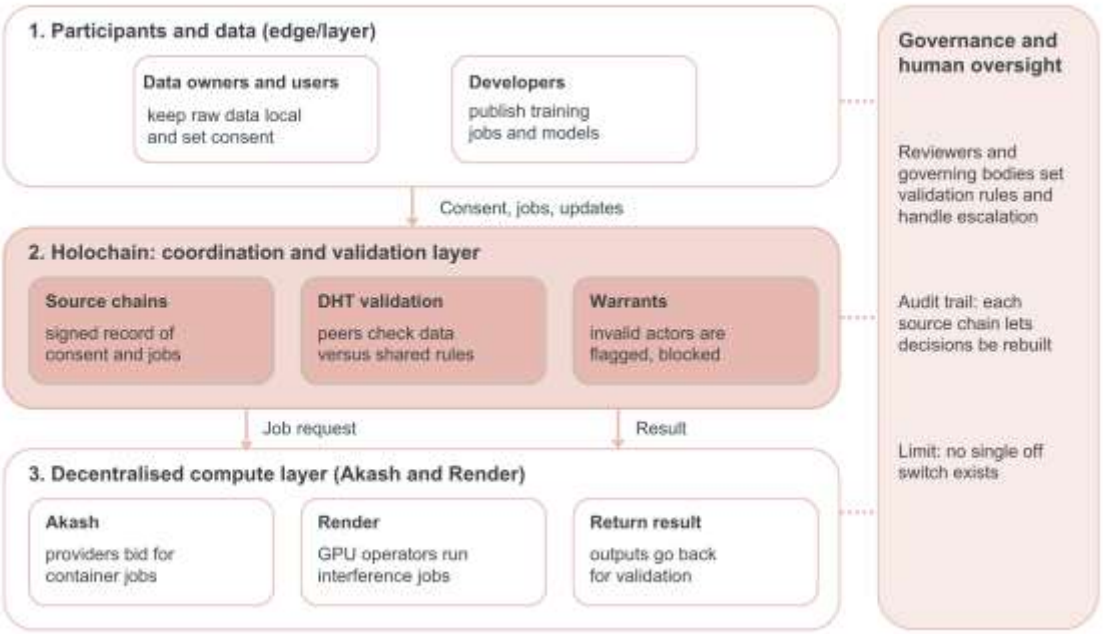


Figure 2 : proposed layered architecture

7. Analysis and Discussion

7.1 How the framework answers the problems of centralised AI

Problem (from the literature)	Mechanism in the framework	Limitation
Pooled data can be breached or misused (Carlini et al., 2021; Shokri et al., 2017)	Raw data persists on owner's source chain; only consented items or model updates are shared (McMahan et al., 2017)	Model updates and outputs can still leak information about training data
Loss of control over data	Consent and usage terms logged as signed entries and peer audited (Holochain, n.d.-b)	Revoking data already used in training is difficult
Concentrated power and vendor lock in (Opara Martins et al., 2016; Widder et al., 2023)	Many independent compute providers on two marketplaces; none control the system	Marketplace demand and supply can concentrate (Render Network Foundation, n.d.)
Need for GPUs	Heavy jobs delegated to Akash or Render	Trust, pricing, and availability varies across providers
Trust in results	Peer validation of records and results per agreed rules	Deterministic rules needed for validation, which cannot by themselves judge model quality or safety (Holochain, n.d.-b)
Weak accountability	Signed source chains provide an audit trail; warrants block rule-breakers	Accountability only extends to protocol layer; does not identify real world operator

International Journal for Innovative Engineering and Management Research
PEER REVIEWED OPEN ACCESS INTERNATIONAL JOURNAL
www.ijiemr.org

7.2 Workflow of a decentralised AI job:

Figure 3 follows one job through the framework. A developer advertises a job specification and a data request. Data owners determine consent. Each consent decision is a signed entry. The request is validated by peers against the shared rules and then a coordinating agent gets compute from Akash or Render. The provider executes the job, the result is validated by peers and then accepted and recorded with provenance. Reviewers can review any stage and escalate.

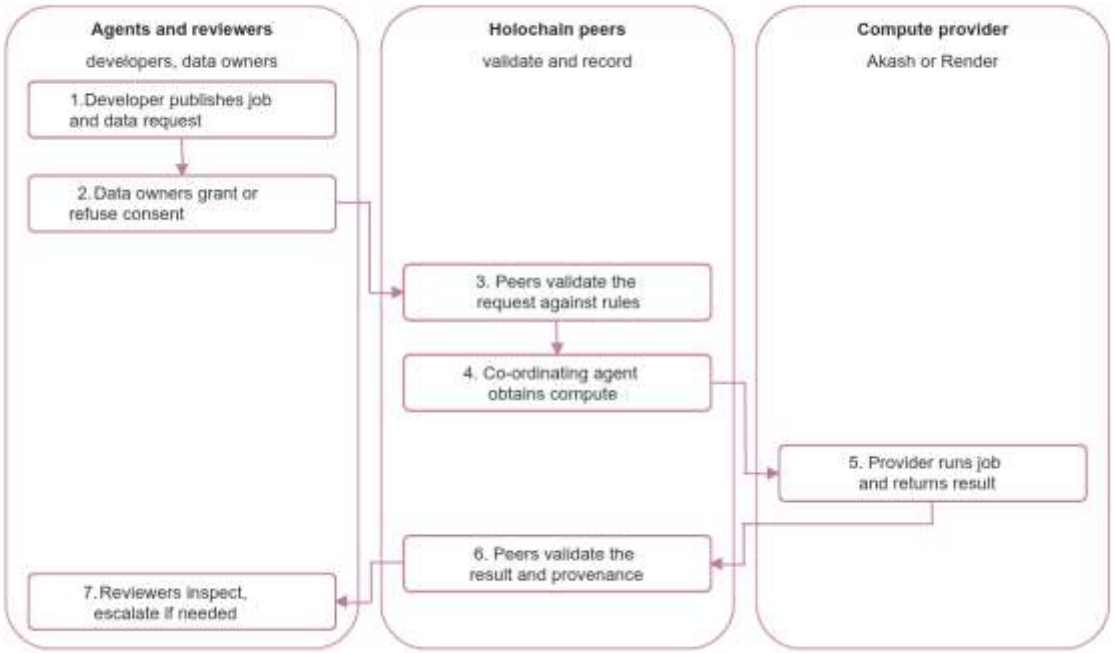


Figure 3: workflow of one decentralised AI.

7.3 Technical limits: verification, privacy, latency, and coordination

Verification:

Peers can verify records, signatures, and compliance with rules; however, they cannot guarantee that the calculations made on a GPU were done appropriately. There are three ways to check approximately and cost wise from the lowest to the highest one. The first is spot check: a certain part of the work has to be done again. The second method is redundant execution, which means running the same task on two different providers and then comparing the final results. The costs of redundant execution approximately double. The third method is to use

cryptographic techniques (Walfish and Blumberg, 2015) which can prove that the calculation was done, and training was completed (Jia et al, 2021). The last two methods are not so effective due to their high overheads and the questions of their reliability, so currently, only risky jobs should be verified more closely.

Privacy:

Federated learning and secure aggregation keep raw input as well as individualized adjustments at their native locations (Bonawitz et al., 2017; McMahan et al., 2017). However, when outsourcing computations this practice changes the situation: anything that is sent to someone else is already available to that someone else. Users can manually opt for what should be transferred, while running the job in the hardware protected environment. None of these solutions can eliminate the possibility of getting the training data from the model per se, and thus the sovereignty claims should be limited by technical possibilities.

Latency:

Publishing data in a Distributed Hash Table, verifying through peers, and transaction processing in a cloud market consumes time. This is fine when it is about batches that may take several minutes or even hours, but it does not work well for providing services for real time queries. One possible strategy is to cash a service provider before the request and to utilize Holochain only for granting access and auditing without wasting time for market negotiations.

8. Major Findings:

1. Decentralisation mitigates but does not eliminate the privacy and concentration issues of centralised AI. Storing raw data with the data owner and distributing compute to the providers solves the risks associated with pooled data and lock in, but still provides opportunities for leaking information (Carlini et al., 2021; Shokri et al., 2017).
2. Holochain is to be utilised only as the coordination, consent, and provenance layer. The agent oriented architecture scales well with the increase in participants and lacks the need for global consensus but was not designed for GPU computations and, thus, cannot solve all the issues (Holochain, n.d.-a).
3. Outsourcing computation by incorporating Akash and Render is essential as it raises the compute power of the system compared to a stand alone holochain system.

4. Ensuring the validity of outsourced computations is an open problem in technology. Peer verification can verify records and policies, but not whether a GPU task has been performed correctly; other solutions are either expensive or under investigation (Jia et al.,2021; Walfish & Blumberg, 2015).
5. Akash and Render are complements. While Akash is appropriate for container based serving and other tasks, Render provides large amounts of consumer GPUs.
6. The evidence base is limited. Literature on the use of Holochain in AI applications is limited to peer reviewed articles, while information on the Akash and Render platforms is available mainly from the project materials and, therefore, conceptual in nature.

9. Suggestions and Recommendations:

9.1 For system design:

NO.	Recommendation	Reason
1.	Akash and Render to be put behind a single interface so work can move between them	Lock in and price risk is significantly reduced
2.	For Interactive serving pre lease capacity; Use holochain specifically for access control and audit only	Helps prevent lag and market delay on the request path
3.	Personal and classified data should not be stored in shared entries, store reference and consent separately	Avoids conflict with deletion rights
4.	Specify reviewer responsibilities, escalation process and shutdown process prior to deployment	Central off switches are absent.

9.2 For Governance Policy:

- Publish the validation rules and the process of updating them, so that everyone will know who has the ability to change the system and how.
- Have multiple overlapping decision centres such as developers, independent validators and data owners' representatives, rather than one council, consistent with polycentric governance (Ostrom, 2010).
- Check for compliance with data protection obligations, specifically regarding how consent, access and erasure requests are addressed when the entries are append-only (European Parliament & Council of the European Union, 2016).

10. Conclusion:

The paper explores whether a decentralised AI system built on Holochain, incorporating Akash and Render to provide GPU computing, could solve privacy, sovereignty and security concerns associated with a centralized AI system. The quantitative analysis reports a positive result for

the above proposal, due to its agent centric nature Holochain becomes well suited for those components of AI systems that involve issues of ownership, consent, provenance, and validation, which ensure that data belongs to their owners and scale without global consensus. However, holochain cannot provide the GPUs needed for training and inference, and this issue is solved by the compute marketplaces mentioned above.

The paper provides an architectural framework with layers. A key take away of the paper is that distributing authority increases accountability and prevents from having a monopoly on power, and also eliminates a single point at which one can shut down a system. That is why human oversight and escalation of authority needs to be embedded into the system.

However, the study has clear limitations. It is theoretical. In the absence of adequate evidence in relation to Holochain, Akash, Render, the framework can be considered a potentially useful design for future research but not as an established alternative to centralised AI.

References:

1. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*, 1175–1191.
2. Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., & Raffel, C. (2021). Extracting training data from large language models. *Proceedings of the 30th USENIX Security Symposium*, 2633–2650.
3. de Vries, A. (2018). Bitcoin's growing energy problem. *Joule, 2*(5), 801–805.
4. Dinh, T. N., & Thai, M. T. (2018). AI and blockchain: A disruptive integration. *Computer, 51*(9), 48–53.
5. European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union, L 119*, 1–88.
6. Harris-Braun, E., Luck, N., & Brock, A. (2018). *Holochain: Scalable agent-centric distributed computing* [Technical white paper]. Holochain. https://github.com/holochain/holochain-proto/blob/whitepaper/holochain.pdf

7. Holochain. (n.d.-a). *Holochain developer documentation*. https://developer.holochain.org
8. Holochain. (n.d.-b). *Holochain developer documentation: Validation and warrants*. https://developer.holochain.org
9. Jia, H., Yaghini, M., Choquette-Choo, C. A., Dullerud, N., Thudi, A., Chandrasekaran, V., & Papernot, N. (2021). Proof-of-learning: Definitions and practice. *2021 IEEE Symposium on Security and Privacy (SP)*, 1039–1056.
10. Akash Network. (n.d.). *Akash Network documentation*. https://akash.network/docs
11. Akash Network. (2018). *Akash Network whitepaper*. https://akash.network
12. Keršič, V., & Turkanović, M. (2026). *A Review on Building Blocks of Decentralized Artificial Intelligence.*
13. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2019). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine, 37*(3), 50–60.
14. Lu, Y., Huang, X., Dai, Y., Maharjan, S., & Zhang, Y. (2020). Blockchain and federated learning for privacy-preserved data sharing in industrial IoT. *IEEE Transactions on Industrial Informatics, 16*(6), 4177–4186.
15. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 54*, 1273–1282.
16. Montes, G. A., & Goertzel, B. (2019). Distributed, decentralized, and democratized artificial intelligence. *Technological Forecasting and Social Change, 141*, 354–358.
17. Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. https://bitcoin.org/bitcoin.pdf
18. Opara-Martins, J., Sahandi, R., & Tian, F. (2016). Critical analysis of vendor lock-in and its impact on cloud computing migration: A business perspective. *Journal of Cloud Computing, 5*, Article 4.
19. Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review, 100*(3), 641–672.
20. Render Network. (2020). *Render Network whitepaper*. https://rendernetwork.com
21. Render Network Foundation. (n.d.). Render Network. Retrieved October 24, 2026, from https://rendernetwork.com

22. Salah, K., Rehman, M. H. U., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access, 7*, 10127–10149.
23. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18.
24. Tramèr, F., Zhang, F., Juels, A., Reiter, M. K., & Ristenpart, T. (2016). Stealing machine learning models via prediction APIs. *Proceedings of the 25th USENIX Security Symposium*, 601–618.
25. Vipra, J., & West, S. M. (2023). *Computational power and AI*. AI Now Institute.
26. Walfish, M., & Blumberg, A. J. (2015). Verifying computations without reexecuting them. *Communications of the ACM, 58*(2), 74–84.
27. Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly, 26*(2), xiii–xxiii.
28. Weng, J., Weng, J., Zhang, J., Li, M., Zhang, Y., & Luo, W. (2021). DeepChain: Auditable and privacy-preserving deep learning with blockchain-based incentive. *IEEE Transactions on Dependable and Secure Computing, 18*(5), 2438–2455.
29. Whittaker, M. (2021). The steep cost of capture. *Interactions, 28*(6), 50–55.
30. Widder, D. G., West, S., & Whittaker, M. (2023). *Open (for business): Big tech, concentrated power, and the political economy of open AI*. SSRN. https://doi.org/10.2139/ssrn.4543807
31. Wüst, K., & Gervais, A. (2018). Do you need a blockchain? *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 45–54.
32. Zheng, Z., Xie, S., Dai, H.-N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services, 14*(4), 352–375.