



The intersection of AI and public blockchains is producing a wave of novel use cases that go far beyond speculative hype, especially as technical primitives on both sides mature.

APR 17, 2025



For a deep-tech audience, the most compelling use cases generally leverage one or more of the following characteristics:

- **AI as an agent** (autonomous, value-seeking actor)
- **Blockchain as a coordination substrate** (trustless computation, incentive design, provenance)
- **Shared data and model governance** (decentralized control over high-value AI resources)
- **Verifiable inference and training** (zero-knowledge, MPC, or TEEs)

Here are the most technically intriguing use cases:

Decentralized AI Agents with On-Chain Incentives

What it is: Autonomous AI agents operating on-chain or interacting with blockchain-based protocols to earn, spend, or manage cryptoassets. These can be personal agents (e.g., AI traders, yield optimizers) or network agents (e.g., oracles or validators).

Why it matters technically:

- Requires mechanisms for on-chain reputation, identity, and staking
- Uses reinforcement learning in economic games (e.g., MEV extraction, DeFi arbitrage)
- Pushes the limits of composability in smart contracts with autonomous logic

Example Projects:

- **Fetch.ai** — economic agents for DeFi and logistics
- **Autonolas** — multi-agent systems for DAO coordination
- **AgentLayer (recent)** — deploying AI agents as smart contract services

Decentralized Model Training and Ownership

What it is: Training machine learning models across distributed nodes where model weights, training contributions, and usage are tracked and monetized on-chain.

Why it matters technically:

- Requires secure aggregation (FL, MPC, or TEEs) across untrusted parties
- Verifiable attribution of training contributions for tokenized reward distribution
- Models as NFTs (e.g., fine-tuned LLM checkpoints as tradeable digital assets)

Example Projects:

- **Bittensor (TAO)** — incentivized decentralized ML network
- **Gensyn** — marketplace for distributed training compute
- **Numeraire** — hedge-fund model built on staking ML predictions

On-Chain Provenance and Auditability of AI Models

What it is: Using public blockchains to anchor training data hashes, model parameters, and inference results for full transparency and traceability.

Why it matters technically:

- Enables compliance with AI governance frameworks (e.g., EU AI Act, HIPAA)
- Facilitates reproducibility in science and medicine (e.g., AI-based diagnostics)
- Introduces the concept of *AI model lineage* as a public asset

Example Projects:

- Ocean Protocol — tokenized data marketplace with provenance tracking
- OpenMined / Gaia-X — data federations for regulated AI

Zero-Knowledge ML (zkML)

What it is: Using zk-SNARKs or zk-STARKs to prove that an ML inference was performed correctly, without revealing the model weights or input data.

Why it matters technically:

- Verifiable inference allows use of proprietary models in trustless settings
- Protects sensitive inputs (e.g., in healthcare or identity verification)
- Enables trustless AI agents in on-chain games, voting, or escrow

Example Projects:

- Modulus Labs — zkML for trustless AI games and DeFi
- RISC Zero — general-purpose zkVM running ML workloads
- ZKonduit / EZKL — proving circuits for neural networks

AI-Governed DAOs and Autonomous Oracles

What it is: DAOs increasingly use AI to evaluate proposals, summarize debates, even vote. Autonomous oracles can also use LLMs to interpret off-chain data and

submit structured information on-chain.

Why it matters technically:

- Governance bottlenecks can be mitigated by AI-mediated summarization and prioritization
- AI as a subjective oracle layer opens up new markets (e.g., litigation, sentiment analysis)
- Merges natural language understanding with on-chain action

Example Projects:

- Kleros — dispute resolution with potential LLM arbitrators
- Delphi Systems — agent-based proposal evaluation in DAOs

Synthetic Data Markets with On-Chain Incentives

What it is: Token-incentivized creation, curation, and usage of synthetic datasets for training models when real data is scarce or sensitive (e.g., in healthcare, finance).

Why it matters technically:

- Combines generative AI with token economics to address data scarcity
- Smart contracts enforce licensing, quality scoring, and payment for data use
- Enables federated learning across data silos in regulated sectors

Example Projects:

- Synapse AI / Alethea — data monetization via tokenized synthetic content
- Arkhn (early-stage in healthcare) — bridging EHR data to public AI pipelines with provenance

If you're exploring use cases in **healthcare**, I'd call out **verifiable inference for clinical decision support** and **token-incentivized medical data markets** (e.g., synthetic MRI datasets with privacy-preserving training) as particularly relevant.



2 Likes • 1 Restack

← Previous

Next →

Discussion about this post

Comments

Restacks



Write a comment...
