ResonanceNet

A Peer-to-Peer Intelligence Currency

Kristian Pilatovich

pilatovichkristian2@gmail.com

March 2026

Abstract

A purely peer-to-peer intelligence system would allow decentralized artificial intelligence to be trained, improved, and distributed without relying on any central authority. We propose a network where computational work is not wasted on arbitrary hash puzzles but instead directed toward training a shared neural network. Miners compete by improving model performance on a common validation benchmark; blocks are accepted only when validation loss decreases. The result is a blockchain whose security grows alongside a continuously improving AI model — a living, permissionless intelligence that belongs to no one and serves everyone. We introduce Proof-of-Training (PoT), a novel consensus mechanism where the chain's integrity is maintained by the same computation that makes the model smarter. Combined with ResonanceNet, an O(1)-state neural architecture achieving 500–1000x data efficiency over conventional transformers, this system creates the first trustless, self-improving artificial intelligence. The model runs natively on mobile devices, requires no data centers, and is secured by Keccak-256d cryptography with Ed25519 signatures. Lightning payment channels enable instant transactions from launch. Wallet recovery is mandatory, eliminating permanent coin loss.

1. Introduction

The current artificial intelligence landscape is defined by a paradox. Hundreds of billions of dollars flow into AI infrastructure, yet 95% of organizations report zero measurable return [1]. The largest language models require data centers costing $100M+ to train, consume megawatts of power, and remain locked behind corporate APIs. The intelligence they produce is proprietary, mutable, and subject to censorship.

Meanwhile, the cryptocurrency ecosystem has demonstrated that decentralized networks can create and secure immense value — Bitcoin alone exceeds $2 trillion in market capitalization — but the computational work securing these networks produces nothing beyond the security itself. Proof-of-Work mining is, by design, useful only for its own verification.

We propose a synthesis: a network where mining is training. Every block added to the chain represents a measurable improvement in a shared neural network. The work that secures the chain is the same work that makes the AI smarter. No computation is wasted. No central authority controls the model. The intelligence is permanent, immutable, and freely accessible — secured by the same cryptographic guarantees that protect Bitcoin.

This is not an incremental improvement to either field. It is a paradigm shift: the creation of an intelligence currency, where the token's value is intrinsically tied to the quality of a real, verifiable AI model recorded on-chain.

1.1 The AI Crisis

The artificial intelligence bubble of 2024–2026 mirrors the dot-com era in structure but exceeds it in scale:

Metric	Value	Source
Enterprise GenAI investment with near-zero ROI	$30–40 billion	MIT Nanda Lab, 2025 [1]
OpenAI projected operating losses (2028)	$74 billion	Wall Street Journal, 2026
OpenAI cumulative losses (2024–2029)	$140 billion	Deutsche Bank analysis
AI capex spending by hyperscalers (2026)	$527 billion	Goldman Sachs [8]
Potential equity destruction if AI bubble deflates	$33 trillion	Oliver Wyman [2]
DeepSeek incident — market cap erased in one day	$600 billion	January 2025
Total AI sector market capitalization	$41 trillion	CompaniesMarketCap, Mar 2026
NVIDIA market cap alone	$4.6 trillion	February 2026

Table 1. The AI investment crisis in numbers.

The fundamental problem is not artificial intelligence itself but the centralized, capital-intensive paradigm through which it is developed. A single company (OpenAI) projects $74 billion in operating losses for 2028 alone, while its models remain locked behind APIs subject to censorship and rate limiting. ResonanceNet addresses this directly by distributing the cost of training across a global network of miners, each incentivized by token rewards.

1.2 Existing Approaches and Their Failures

Bittensor (TAO) attempts decentralized AI through a network of validators that subjectively score model outputs. This introduces human judgment into consensus — a fundamental weakness. Validators can collude, scores are non-deterministic, and there is no single, objectively improving model. Furthermore, Bittensor uses standard transformer architectures requiring data center-class hardware for inference, negating the promise of decentralization.

Qubic uses "Useful Proof of Work" to train arbitrary neural networks, but produces no single coherent model. Federated learning distributes training but requires a trusted aggregator. Proof-of-Useful-Work proposals (Primecoin, Chia) redirect mining toward mathematical problems but produce no reusable computational artifact.

ResonanceNet is the first system where consensus, security, and intelligence production are unified in a single, objectively verifiable process.

1.3 The Decentralized AI Opportunity

Centralized AI commands approximately $12 trillion in enterprise value, while decentralized AI protocols are valued at roughly $12 billion — a 1000:1 ratio that represents an unprecedented investment opportunity [9]. The blockchain AI market is projected to grow from $6 billion in 2024 to over $50 billion by 2030 at a 42.4% CAGR, with analysts suggesting these figures significantly underestimate actual growth. Bitcoin miners are already pivoting GPU infrastructure toward AI compute, creating natural demand for a protocol that unifies both activities.

2. ResonanceNet Neural Architecture

The neural architecture underlying the network must satisfy constraints that no existing model meets simultaneously: (1) O(1) state per token, enabling inference on mobile devices; (2) deterministic forward pass for independent verification; (3) continuous modular growth as the chain extends; (4) extreme data efficiency to minimize the computational barrier to mining.

2.1 Model Design

ResonanceNet V5 is a recurrent neural architecture that replaces self-attention with three complementary mechanisms. Per-layer pipeline:

```
Input → RMSNorm → Multi-Scale Causal Convolution → +residual → RMSNorm
→ MinGRU (parallel scan) → +residual → RMSNorm → Slot Memory
(cross-attention) → +residual → RMSNorm → SwiGLU FFN → +residual
```

Output: RMSNorm → weight-tied embedding transpose → logits.

2.1.1 Multi-Scale Causal Convolution

Each layer applies N parallel causal convolution branches with kernel sizes K = {3, 7, 15, 31, 63}. This captures local patterns at multiple time scales without the quadratic cost of self-attention.

2.1.2 MinGRU with Log-Domain Parallel Scan

The recurrent component uses Minimal Gated Recurrent Units [3] computed via the Blelloch parallel prefix sum algorithm in log-space. This achieves O(n) training via parallel scan (vs. $O(n^2)$ for attention), O(1) inference per token, and numerical stability through log-domain computation. The recurrence:

$$h_t = (1 - z_t) * h_{t-1} + z_t * \text{h\_tilde}_t$$

2.1.3 Slot Memory

A fixed set of S = 64 learnable key-value slots provides global context through cross-attention. This replaces self-attention's $O(n^2)$ sequence-to-sequence computation with O(n·S) sequence-to-slots computation, where S is fixed. The slots learn to represent persistent concepts — a compressed world model that accumulates knowledge across training.

2.2 Data Efficiency

In controlled experiments, ResonanceNet achieves validation perplexity of 189.37 after training on 86 million tokens at 33.62M parameters. Comparable transformer architectures require approximately 33.5 billion tokens to reach equivalent perplexity at similar scale — a conservatively estimated 388× data efficiency advantage.

Critically, data deficit was observed at approximately 43 million tokens (halfway through training), indicating the model had already extracted maximum information from the available corpus. The true efficiency ratio is therefore closer to 33.5B / 43M ≈ 780×, with theoretical upper bounds of 500–1000× depending on data quality and domain diversity.

2.3 Inference Characteristics

Property	Transformer	ResonanceNet
State per token	O(n) KV cache	O(1) hidden state
Inference compute	O(n^2) attention	O(1) recurrence
Memory at 100K context	~8 GB (7B model)	~2 MB
Memory at 3B params	~6 GB KV cache	~96 MB state
Mobile inference	Impractical	Native
Inference speed (34M, phone)	N/A	2,000–8,000 tok/s
Inference speed (34M, GPU)	N/A	50,000–100,000 tok/s

Table 2. Inference comparison: Transformer vs. ResonanceNet.

The O(1) state property is the key differentiator. Any device that can download the tip checkpoint can run the full model. The AI is not a service — it is a possession.

3. Proof-of-Training (PoT)

3.1 Consensus Rule

The consensus rule is a single mathematical inequality:

```
val_loss(checkpoint_new, D_val) < val_loss(checkpoint_parent, D_val)
```

A block is valid if and only if the model checkpoint it contains achieves strictly lower validation loss than its parent block's checkpoint. This rule is objective (loss is a deterministic scalar), verifiable (any node can recompute), and unforgeable (no shortcut to reducing loss other than genuine training).

3.2 Mining Process

1. Miner downloads the tip checkpoint. 2. Miner selects ANY training data of their choice. 3. Miner trains for N steps. 4. Miner evaluates on canonical validation set D_val. 5. If val_loss improved: construct block, sign, broadcast. 6. Network verifies by recomputing val_loss. 7. Block accepted if loss matches within tolerance (2%). 8. Miner receives reward.

The miner's choice of training data is unconstrained. This creates a natural data marketplace: as the model exhausts public data, miners must acquire or generate novel data. The token (RNT) becomes the currency of this data economy. A single miner with unique, high-quality data and a consumer GPU (e.g., RTX 5090) can outperform an entire data center training on exhausted public datasets.

3.3 Validation Protocol

Full validation requires a forward pass over the validation dataset — approximately 10,000× cheaper than training. Light validation (header-only) checks structural rules without GPU: height increments, hash chain, loss monotonicity, growth field correctness, signature validity. This enables lightweight nodes on consumer hardware.

3.4 Security Analysis

Attack Vector	Defense
Fake loss reporting	Checkpoint hash in header; any validator recomputes loss
Overfitting to D_val	Diminishing returns; governance D_val rotation via supermajority
Miner collusion	Mining is competitive; only the best improvement wins the block
51% attack	Requires majority GPU power AND superior training data
Long-range attack	Cumulative score makes deep reorgs extremely expensive
Chain stall (no improvement)	Continuous growth mechanism creates new model capacity

Table 3. Security analysis of Proof-of-Training.

4. Continuous Model Growth

Unlike fixed-schedule growth (e.g., +1 layer every 200 blocks), ResonanceNet employs continuous, adaptive growth. The model architecture expands with every block, and growth rate accelerates when loss improvement stagnates — creating organic difficulty regulation.

4.1 Growth Algorithm (Consensus-Critical)

Every node must compute identical growth for a given block. If loss improved: stagnation resets to 0, d_model += BASE_GROWTH (2). If loss did not improve: stagnation increments, d_model += BASE_GROWTH * (1 + stagnation / PATIENCE). Layers are added when cumulative d_model growth crosses LAYER_THRESHOLD (128). Constants: MAX_D_MODEL = 4096, MAX_LAYERS = 48.

4.2 Growth Trajectory

Block	d_model	Layers	Params (approx)	Equivalent Transformer
0 (genesis)	384	6	33.6M	—
100	~584	7	~65M	—
500	~1,384	13	~500M	GPT-2
1,000	~2,200	20	~2B	—
2,000	~3,500	30	~10B	LLaMA-7B (at 500–1000x eff.)
3,000+	4,096 (cap)	48 (cap)	~30B	LLaMA-13B+ (at 500–1000x eff.)

Table 4. Model growth trajectory with continuous growth policy.

4.3 Organic Difficulty

There is no numeric difficulty target. Difficulty is emergent: when loss improves often, growth is slow and blocks are frequent. When loss plateaus, growth accelerates, new parameters create room for improvement, and blocks resume. This creates natural cycles without any difficulty adjustment algorithm — a self-regulating system driven entirely by the mathematics of learning.

5. Cryptographic Foundation

All hashing uses Keccak-256d (double Keccak-256), the original Keccak submission with 0x01 padding — identical to Ethereum's keccak256. Double hashing protects against length extension attacks. Digital signatures use Ed25519 (128-bit security, 64-byte signatures, deterministic signing, fast batch verification). P2P transport is encrypted with ChaCha20-Poly1305 AEAD.

6. Tokenomics

Parameter	Value
Ticker	RNT
Total supply	21,000,000 RNT
Base block reward	50 RNT
Halving mechanism	Adaptive (based on effective supply, not block count)
Minimum denomination	1 resonance = 10■■ RNT
Performance bonus	Up to 2× base reward for larger loss improvements
Bech32 address prefix	rnt1 (mainnet), trnt1 (testnet)

6.1 Adaptive Emission

Unlike Bitcoin’s fixed halving schedule (every 210,000 blocks), ResonanceNet halves the block reward based on effective circulating supply. Coins locked in wallets past their heartbeat deadline or proven lost are subtracted from effective supply. This means if many coins are lost, halving happens later, maintaining healthy emission to replace lost supply. The emission rate naturally adjusts to the real economy rather than an arbitrary block count.

6.2 Value Formation

Phase 1 — Free data era: public datasets provide easy improvements, miners earn quickly. Phase 2 — Data scarcity: miners must acquire novel data, establishing a price floor. Phase 3 — Data marketplace: RNT becomes the medium of exchange for training data. Unlike speculative tokens, RNT’s value is anchored to the cost of improving a real AI model.

7. Mandatory Wallet Recovery

An estimated 20% of Bitcoin supply is permanently lost. ResonanceNet eliminates this through mandatory recovery policies set at wallet creation. No wallet can be created without selecting one of three recovery strategies:

Heartbeat Recovery: Owner sets an interval (10,000–500,000 blocks) and a recovery address. A periodic heartbeat transaction proves continued access. If missed, the recovery address can claim funds after a timeout period.

Social Recovery (M-of-N): Owner designates N guardian public keys and threshold M. If access is lost, M guardians sign a recovery transaction. Owner can cancel within a waiting period (1,000+ blocks).

Emission Return: After a long inactivity period (200,000+ blocks), funds return to the emission pool as future block rewards, recycling lost coins back into the economy.

Encrypted backups use Argon2id key derivation (64MB memory-hard) + AES-256-CBC. The resulting encrypted file can be stored anywhere (cloud, USB, printed QR code) without revealing wallet contents. BIP39 seed phrases (24 words) provide the master backup.

8. Lightning Payment Channels

ResonanceNet blocks are produced at variable intervals (training time = minutes to hours). Lightning payment channels are included from v0.1, enabling instant micropayments from day one.

Use cases: streaming micropayments for training data (pay per MB), inference API billing (pay per token), mining pool payouts, and general peer-to-peer payments. The implementation follows the BOLT specification with Sphinx onion routing for multi-hop privacy, HTLC (Hash Time-Locked Contracts) using Keccak-256d, and BOLT11-style invoices with “rnt” prefix.

9. Multi-GPU Backend

ResonanceNet supports three GPU backends through an abstract GpuBackend interface: CUDA (NVIDIA), Metal (Apple Silicon), and Vulkan Compute (AMD, Intel, Qualcomm, Mali). Training and inference code never call GPU APIs directly — all computation goes through the abstraction layer. Auto-detection selects the best available backend at runtime.

Platform	Backend	Targets
NVIDIA (desktop, server)	CUDA 12.0+	Training + inference
Apple Silicon (Mac, iPhone, iPad)	Metal	Training + inference
AMD, Intel, Qualcomm, ARM Mali	Vulkan Compute 1.3+	Training + inference
Android phones (2024+)	Vulkan	Inference: 2–8K tok/s at 34M

Platform	Backend	Targets
iPhone (2024+)	Metal	Inference: 3–8K tok/s at 34M

Table 5. GPU backend support matrix.

10. Network Protocol

Nodes communicate over an encrypted P2P network with block propagation (headers first), chunked checkpoint transfer (1MB chunks for 65MB+ checkpoints, parallel download from multiple peers), DNS seed peer discovery, and native SOCKS5 proxy support for Tor/I2P. New nodes download headers, verify the chain, and fetch only the tip checkpoint. The wire protocol uses RNET magic bytes (0x52 0x4E 0x45 0x54) with Keccak-256d checksums.

Fork resolution selects the block with lowest val_loss at the tip. For deep forks, the chain with lowest cumulative val_loss from the fork point is selected. Finality is achieved after 6 confirmations.

11. Comparison

Property	Bitcoin	Bittensor (TAO)	ResonanceNet
Consensus	PoW (SHA-256d)	PoS + subjective	PoT (val_loss)
Mining output	Heat	Model marketplace	Single improving AI
Verification	Hash < target	Validator vote	Forward pass
Useful compute	0%	Partial	100%
Mobile inference	—	No (transformer)	Yes (O(1) state)
Objectivity	Full	Subjective	Full
Data efficiency	N/A	1x (transformer)	500–1000x
Inference memory (3B)	N/A	~6 GB KV cache	~96 MB
Censorship resistance	Financial only	Partial	Full (AI + financial)
Wallet recovery	Optional	Optional	Mandatory
Lightning payments	Via LN (added 2018)	No	Native from v0.1

Table 6. Protocol comparison.

12. Roadmap

Phase	Description
Phase 0 — Genesis	Architecture proven. 33.6M model, 500–1000x efficiency. V4 codebase complete (~80–100K LOC). Testnet o
Phase 1 — Mainnet	Genesis block. Public mining. CLI wallet with integrated mining. Lightning channels. Seed nodes operational.
Phase 2 — Growth	Model surpasses 1B parameters. Data marketplace emerges. Mobile inference engine released. Community G
Phase 3 — Ecosystem	Third-party applications. Multi-modal extensions. The model becomes global AI infrastructure.

13. Network Constants

Parameter	Mainnet	Testnet
Magic bytes	0x524E4554 (RNET)	0x544E4554 (TNET)
P2P port	9555	19555

Parameter	Mainnet	Testnet
RPC port	9554	19554
Lightning port	9556	19556
Bech32 prefix	rnt1	trnt1
Base58 P2PKH prefix	R	T
BIP44 coin type	9555	9555

14. Conclusion

ResonanceNet is not another cryptocurrency. It is not another AI project. It is the unification of both into something that has not existed before: a peer-to-peer intelligence currency.

Every block mined makes the network's AI measurably smarter. Every token earned represents real computational work that produced real intelligence. The model grows continuously, verifiably, and permanently — secured by cryptography, driven by economics, and owned by no one.

The AI industry's current trajectory — centralized, capital-intensive, producing diminishing returns — is unsustainable. OpenAI projects $74 billion in losses for a single year. The total AI sector at $41 trillion in market capitalization rests on infrastructure that a single disruption (DeepSeek, January 2025) can erase $600 billion from overnight.

ResonanceNet offers the alternative: a decentralized intelligence that improves continuously through the competitive efforts of a global network of miners, each incentivized to find better data and better training strategies. The model runs on any device — from data centers to mobile phones. No API keys. No censorship. No single point of failure.

"The model is the chain. The chain is the model. The intelligence is permanent."

References

[1] Nanda Lab, MIT Media Lab. “Enterprise AI Investment Returns.” August 2025.

[2] Oliver Wyman. “How An AI Bubble Burst Could Shake Global Financial Markets.” January 2026.

[3] Feng, L., Tung, F., Ahmed, M.O., Bengio, Y., & Hajimirsadeghi, H. “Were RNNs All We Needed?” 2024.

[4] Blelloch, G.E. “Prefix Sums and Their Applications.” Carnegie Mellon University, 1990.

[5] Bertoni, G., Daemen, J., Peeters, M., & Van Assche, G. “The Keccak Reference.” 2011.

[6] Nakamoto, S. “Bitcoin: A Peer-to-Peer Electronic Cash System.” 2008.

[7] Bernstein, D.J., et al. “High-speed high-security signatures.” Ed25519, 2012.

[8] Goldman Sachs Research. “Why AI Companies May Invest More than $500 Billion in 2026.” December 2025.

[9] CoinDesk. “How Decentralized AI Is Leveling the Playing Field.” February 2026.

[10] Wall Street Journal. “OpenAI Financial Projections.” March 2026.

[11] CompaniesMarketCap. “Largest AI Companies by Market Capitalization.” March 2026.

ResonanceNet — Peer-to-Peer Intelligence Currency

pilatovichkristian2@gmail.com