
Shannon: Compression as Capital, A Coin for the Intelligence Age

Liam Zebedee

Abstract

The utility of a store of wealth depends on its properties. Bitcoin is useful insofar it is unique - it has a fixed supply, it is uncensorable, and completely digital. This makes it a useful complement to gold and fiat currency.

Hutter proposed that intelligence is equivalent to compression. The better we can predict the world, the fewer bits it takes to describe it. Thus, as AI gets better, compression gets better.

Compression is interesting as compression gains are objective and easy to verify, like Bitcoin's proof-of-work. If you can make a .zip file 10% smaller, it is easy to verify.

We hypothesise a new Bitcoin-like network where the "work" is provable improvement in data compression. Miners compete to compress data more efficiently, which is a direct proxy for improvements in artificial intelligence.

The result is a new form of money whose underlying value tracks advances in intelligence itself.

1 Introduction

Hutter's [seminal idea](#) that AI equals prediction equals compression suggests that the better a system can predict data, the better it can compress it. To grasp the global utility of this concept, consider that BitTorrent accounts for approximately 4% of global internet traffic. With global internet traffic exceeding 33 exabytes daily, 4% amounts to roughly 0.528 exabytes per day, or about 192 exabytes per year. Given the average cost of transmitting data at \$2.59 per gigabyte, this translates to about \$500 billion annually. A 0.1% improvement in compression efficiency could save roughly \$500 million per year. A decentralized system that enables anyone to provably improve compression efficiency can drive continuous global cost savings and innovation.

todo: figures based on public datasets, roughly correct for now

2 Background

In the 2020s, advancements in AI infrastructure have dramatically improved predictive capabilities, with intelligence roughly doubling every 12 months (Altman, 2023). While much of this progress has been directed toward generative AI applications, the potential for leveraging these advancements in data compression has been relatively overlooked. One major barrier to innovation in this area is the complexity and coordination required to integrate new compression standards. Implementing a new video or image codec (like WebM), for instance, involves extensive industry collaboration and significant social coordination costs.

Despite these challenges, the rise of open-source communities and techniques like quantization highlight the opportunity to build a public contest environment, similar to Kaggle, where developers can compete to create superior data compressors using machine learning models. Such a system

would not only foster innovation but also make it possible to continuously improve compression efficiency on a global scale.

A permissionless system that incentivizes compression improvements—effectively a proxy for intelligence—could align economic incentives with the ongoing AI boom. Such a coin would effectively represent the value of collective intelligence and compression efficiency, tying its worth to the entire AI market.

3 Compression as a proxy for intelligence

Learning is the process of constructing internal models that capture patterns in the world. The better these models, the more accurately they can predict what comes next—whether in language, vision, or any other domain. Hutter’s theory asserts that this act of finding patterns is the essence of intelligence. If a model can predict the next bit in a stream with high accuracy, it implies it has captured the underlying logic generating that stream. The more accurate the model’s predictions, the less new information each piece of data carries. Fewer surprises mean fewer bits are needed to represent the data. Compression happens when the model knows so much that most of the data becomes predictable—and only the unpredictable parts need to be stored.

4 Incentivising compression is incentivising public intelligence

Open source software is one of the clearest examples of a public good with outsized private impact. Linux, for instance, powers the vast majority of web servers, cloud infrastructure, and smartphones. Though no one owns it, trillions of dollars in value depend on it. Its existence proves that distributed, voluntary coordination can outcompete proprietary systems when the goal is well-defined and the feedback loop is tight. This pattern repeats in AI: LLAMA and its derivatives have accelerated model development worldwide, enabling rapid advances in quantization, prompting, and efficiency—all from a freely available model. These systems are not valuable despite being open; they are valuable because they are open.

Compression, when linked to intelligence, offers the next logical domain for public coordination. If intelligence is the ability to predict—and prediction enables compression—then progress in compression is measurable progress in intelligence. But unlike traditional AI benchmarks, compression is grounded in information theory: it has clear, objective metrics. By creating a system that rewards improvements in compression, we directly incentivize public contributions to collective intelligence.

Bitcoin demonstrated that a decentralized network can create and maintain a globally valuable asset by rewarding a specific kind of work—hash computation. Imagine if that work wasn’t wasted energy but improved intelligence. A token backed by verifiable gains in compression would reward the production of public models that make data smaller, cheaper, and more transmissible. It turns raw intelligence—encoded in compressors—into a public asset. This isn’t just a better benchmark or contest. It’s an economic mechanism for aligning individual effort with global understanding.

todo: refer to this <https://kk.org/thetechnium/public-intelligence/>

5 System Design

The system modifies Bitcoin’s UTXO model to reward compression instead of hashcash. There is a fixed supply of 21 million tokens. Instead of mining by guessing nonces, tokens are earned by submitting models that compress a public dataset more efficiently.

Kolmogorov complexity is the length of the shortest possible program that can generate a given dataset. Formally, the complexity of data D given model class M is:

$$K(D) = \min_M (Size(M) + Size(M(D)))$$

The goal of this network is to minimize $K(D)$: to reward models that reduce the total description length of the dataset. It is a global optimization loop for compressive power.

(??? define this more clearly, concretely, refer to fabrice bellard's https://bellard.org/ts_sms/tinystories)

Nodes propose their compression models as containerised software which uses CPU and GPU. They are scored on a basis of space and time, allowing for models which optimize for fast compression (low time) or small compression (low space).

Models will in 2035 be proposed as ZK proofs, which allow nodes to demonstrate compressive efficiency in $O(\log(N))$ efficiency. Afterwhich, nodes of the network will download them and start mining on the dataset by compressing it.

Block rewards are awarded similar to Bitcoin. The regular proof-of-work challenge is replaced by a compression challenge on a randomly chosen subset of files in the dataset. We expect that due to the immense size of the dataset (PB's in size), compression challenges will incent nodes to collaboratively shard data in exchange for a share of block rewards like mining pools shard compute capacity.

Difficulty, instead of representing hashpower, represents compressive power - which is modulated by changing the length of the compression context. Faster compression means difficulty increases which requires nodes to compress more files.

Tokenomics are designed based on how the network's utility improves with better compression models. Better compressors reduce the total size of the dataset. Smaller dataset → more room to store new data using same physical capacity. Every time someone submits a better model, they increase the system's data capacity. This lets the system sell more storage — creating new token demand.

6 Tokenomics

21 million shannon. An issuance curve mimicking Bitcoin's. Fees paid to the miners for storing data in the dataset. Each model proposer receives royalties based on the compression gains by their algorithm. For example, a 1% increase in compression efficiency converts to a 1% increase in the supply. (??? there's a lot of designs here, not sure which one for now is best)

The issuance schedule of the native asset will likely follow a trajectory similar to Bitcoin—gradually decreasing over time to incentivize early contributions and stabilize long-term value. In Bitcoin, this dynamic has been shaped not only by the halving schedule but also by broader technological and mimetic forces. On the technological side, mining evolved from CPUs to GPUs, then to FPGAs and ASICs, dramatically increasing computational efficiency. In parallel, Bitcoin's adoption scaled from hobbyists to institutions and governments, driven by growing mindshare and memetic appeal. This dual track—advances in physical hardware and symbolic adoption—forms a kind of distributed, social computation around the asset.

AI, and by extension compression will follow similar path. Macro forces are emerging: OpenAI has positioned itself as a proprietary intelligence provider, while Meta's competitively released the series of LLAMA models which seeded a flourishing open-source ecosystem. Software innovation has exploded, as seen in techniques like quantization (e.g., TheDude's GGML) that radically reduce model size and inference cost. Compression performance is being pushed by all these factors—hardware, software, and the growing cultural and institutional investment in machine intelligence.

Q: how do we best incentivise new intelligence/compression models to be released? nodes will perform compression/decompression? should an ongoing proportion of block rewards go towards these inventors? I think so. I don't think it should go to the last best model proposed either. I think it should be gradually better.

7 Conclusion

This system monetizes compression improvements as a public good. It turns AI-driven prediction into capital, creating a flywheel of utility, incentive, and innovation.

References