

Bittensor Double-Down Investment Memo (Aug 2024)

- **Bittensor is a decentralized network that incentivizes the creation of machine intelligence.** Miners aggregate various resources needed to produce machine intelligence (e.g. compute, data, bandwidth, electricity); validators ensure that resources are utilized in a way that drives economic value; and users benefit from an ever-increasing gradient of more powerful and useful AI apps.
- **Bittensor's native token, TAO, is one of EV3's core investments.** We started building a position in Q1'23 when the network consisted of roughly \$30m of GPUs powering intelligence comparable to GPT-2. Bittensor has since become the leading decentralized AI network by [mind share](#), [market cap](#) and [trading volumes](#), connecting an estimated \$150m worth of GPUs and powering intelligence comparable to GPT-3.5—or better in certain domains.
- **We believe decentralized AI's "killer app" will emerge from the Bittensor ecosystem in the next 18 months**, driving a material re-rating of TAO to valuations comparable to leading centralized AI companies (2-15x upside to OpenAI/Anthropic) and smart contract platforms (10-30x upside to Ethereum/Solana). Bittensor's 'network-of-networks' architecture features 50+ subnets, each of which represents a distinct shot on goal for spawning a killer AI app.

Investment Rationale

The investment case for TAO boils down to three beliefs:

1. Bittensor will spawn decentralized AI's first killer app in the next 18 months.
2. Successful apps built in the Bittensor ecosystem accrue value to TAO.
3. Despite its (nominally) high valuation, Bittensor's success is nowhere near priced in.

Further Reading

- Official: [Website](#), [Medium](#), [Whitepaper](#)
- Explorers: [TaoStats](#), [TaoPill](#), [BittensorWiki](#)
- Podcasts: [Novelty Search](#), [Bittensor Hub](#), [OSS Capital](#), [Proof of Coverage](#)
- Research: [Stateless](#), [Collab+Currency](#), [M31](#), [Messari](#), [Delphi](#), [Teng Yan](#), [TensorExchange](#), [DMD](#)

Apps/Demos built on Bittensor

BitMind deepfake image detection Dippy roleplay chatbot GoGo browser agent Corcel chatbot Omega any-to-any Datura chatbot	NicheTensor text-to-image MakeltAQuote text-to-image TensorAlchemy text-to-image Bettensor sports odds Sportstensor sports odds ScorePredict sports odds	GraphiteAI graph simulation SkinScan cancer detection MyShell text-to-speech ItsAI deepfake text detection Gen42 chatbot NextPlace home prices
--	---	---

Can the Bittensor ecosystem spawn a “killer AI app” in the next 18 months?

Today’s state-of-the-art AI lives within the walls of centralized companies and closed-source code. The open-source community relies on the temporary generosity of Meta / Mistral / Stability to keep up, but the performance gap remains wide. Bittensor is the only serious contender incentivizing open-source AI at a scale that can compete with centralized providers: at current prices, TAO emissions represent >\$3m of daily incentives to open-source AI developers. Our diligence focused on determining if TAO emissions are being used to incentivize open-source AI that has a viable path to beating state-of-the-art centralized AI.

TAO emissions are split across 46+ subnets focused on distinct parts of the AI stack. Miners compete to produce the best (i.e., highest-ranked) responses, while minimizing offchain costs. Crypto- incentives push miners to discover ever-more creative optimizations at breakneck speed: many subnets have seen 80%+ performance improvements in the first few months post-launch. There are subnets focused on:

- Data collection ([21](#)), labeling ([33](#)), and storage ([13](#))
- Model pre-training ([9](#)), fine-tuning ([37](#)), inferencing ([4](#)), and evaluations ([15](#))
- Generating images ([26](#)), 3D objects ([17](#)), and protein structures ([25](#))
- Agents for roleplay ([11](#)), productivity ([20](#)), and DeFi ([10](#))
- Trading stocks ([28](#)), crypto ([8](#)), sports ([30](#)), and prediction markets ([6](#))
- Detecting fraudulent text ([32](#)), images ([34](#)), payments ([14](#)), and diseases ([46](#))

Of the 46 active subnets, we estimate 10 are building real products with strong momentum (30% of total emissions), 25 are promising but early (50% of emissions), and the rest are stagnant and/or likely to be replaced as subnet competition intensifies (20% of emissions). Some of our favorites include:

1. SN25, built by [Macrocosmos](#), incentivizes protein folding simulations. The state-of-the-art in protein folding is Google’s [AlphaFold3](#), which is [publicly-accessible](#) but with limits on usage (20 jobs/day), complexity (5k tokens/job), monetization (non-commercial only), and censorship (e.g. viral pathogens). The second-best existing option for protein folding simulations is [ROSIE](#), a hosted solution built by researchers at Johns Hopkins with a [10+ day wait-time](#). SN25 has already done [more folding jobs](#) in its first three months than ROSIE has in the past three years.
2. SN34, built by [BitMind](#), incentivizes deepfake image detection models. BitMind developed a hard-mixture-of-experts methodology that breaks down general deepfake detection into task-specific modules (e.g., face recognition, object detection, frequency analysis). The latest benchmarking of generalized deepfake detection models was [published](#) by Chinese researchers in Oct’23. SN34 miners [outperformed](#) the leading model on every metric, achieving an accuracy of 88% vs 82%. You can query the models with real and/or fake images by uploading them at [BitMindID](#).
3. SN2, built by [Inference Labs](#), verifies the authenticity of models via zk-verified inference. The subnet has attracted [4k+ CPUs and 30TB+ RAM](#) in four months, surpassing Akash and Flux by 2-3x. Several protocols including [Bengi](#) (for yield optimizations) and [TrustMachine](#) (for contract audits) have signaled their intent to use SN2-powered verified inference in production.
4. SN13, built by [Macrocosmos](#), incentivizes the scraping of massive text datasets from online sources. Scraping fresh text data from across the internet plays a critical role in training the next generation of LLMs. The most notable crypto competitor, [Grass](#), has 2m users and a ten-figure [private valuation](#). Grass released its first dataset of [550m Reddit posts](#) in July. In August, Macrocosmos released its first dataset: the top miners each scraped >300m posts, and the top 10 miners scraped [1.3B Reddit and Twitter posts](#) in total—surpassing Grass by a factor of 2x.

5. SN11, built by [Impel Labs](#), incentivizes role-playing LLMs for companion apps. The subnet is [already beating GPT-3.5 Turbo](#) on emotional intelligence benchmarks with models that are an order of magnitude smaller (8B vs 100B parameters). Impel Labs runs Dippy, an AI companion app with 280k+ MAUs. In November, Dippy will [replace its closed-source 40B-parameter model](#) with responses from SN11 miners running open-source 13B-parameter models for all its users.
6. SN39, built by [Wombo](#), incentivizes LLMs that run on consumer-grade devices. Currently, the subnet rewards miners who optimize Stable Diffusion XL to run on Nvidia GeForce RTX 4090s. In the past month, the top miner's generation time has gone from [2.3s](#) to [1.7s](#) (26% improvement). The Wombo team built and operates two [image generation mobile apps](#) with >200m downloads which will integrate images generated by SN39 miners.

Momentum	Promising But Early	Likely Zero
Omron (2): zkML verified inference Taoshi (8): time-series prediction Pretraining (9): foundation models Dippy (11): character roleplay LLMs Data Universe (13): data scraping Cortex (18): synthetic text data Vision (19): distributed inference Protein Folding (25): protein folding BitMind (34): deepfake image detection Edgemaxxing (39): consumer-grade AI	Targon (4): low-cost inference Infinite Games (6): prediction markets Subvortex (7): firewalls Horde (12): distributed compute PayPangea (14): payments fraud detection DeVal (15): model evaluations Omega (24): multimodal datasets Omega (21): multimodal any-to-any Foundry (28): S&P500 price prediction ColdInt (29): collaborative training Bettensor (30): sports predictions Sportstensor (41): sports predictions ScorePredict (44): sports predictions NAS (31): neural architecture search It's AI (32): detect AI-generated text Masa (42): structured data LogicNet (35): mathematical analysis BitAgent (20): productivity agents Tensor Alchemy (26): image generation Fine-tuning (37): fine-tuning LLMs Training (38): distributed training Chunking (40): chunking & retrieval Graphite (43): graph optimizations Gen42 (45): coding assistant Safescan (46): cancer detection Three-Gen (17): text-to-3D SocialTensor (23): image generation	MyShell (3): text-to-speech Kaito (5): text embeddings Sturdy (10): DeFi yield optimizer BitAds (16): advertising Smart Scrape (22): sentiment analysis Neural Internet (27): compute ReadyAI (33): text annotation Human Intelligence Primitive (36)

What makes the Bittensor ecosystem special? Bittensor enables any entrepreneur around the world to register a subnet by staking ~\$500k TAO and defining an objective function. Upon registration, a global network of AI developers and infrastructure providers directs their attention and resources towards continuously optimizing towards that objective function. Historically, Bittensor miners are able to achieve 10-25% *monthly* efficiency gains across virtually every layer of the AI infrastructure stack, and new open-source AI models are integrated into subnets within days, if not hours. The registration fee will be reduced over time and the number of subnets will continue to increase, lowering the barrier to joining the network. Entrepreneurs who are world-class at designing useful incentives can - for the first time - build something huge without having to raise capital, hire developers, or scale infrastructure. The speed and breadth of permissionless innovation in Bittensor is unlike any other we've seen since the start of EV3—if Satoshi were (is?) building in crypto today, we think he would be (is?) building on Bittensor.

2. Does the success of Bittensor subnets imply the success of the TAO token?

TAO will fail if successful subnets fork away and/or otherwise disentangle themselves from the network. There's at least eight extractive subnets today with no intentions of driving meaningful long-term value, collecting a few percentage points of total emissions. On a long-enough time horizon even well-meaning subnets will look to capture more of their own value, unless there's a sound economic mechanism that locks them in. That mechanism - [Yuma Consensus](#) - is the core innovation behind Bittensor.

Bittensor's design is best understood through the lens of progressive decentralization. The end-goal has always been consistent: for TAO emissions to incentivize *useful intelligence*, i.e., intelligence that people will *pay for*. In v1, the Opentensor Foundation defined the parameters for 'usefulness'. The [Revolution upgrade](#) in Q4'23 enabled anyone to create subnets with their own *usefulness* parameters; this removes the dependency on the Foundation, but still requires validators to vote on emissions split across subnets. In Q4'24, the upcoming [Dynamic TAO](#) upgrade will remove validator voting and replace it with a market-driven mechanism: subnets will launch their own tokens on top of the Bittensor blockchain, with TAO emissions deposited as single-sided liquidity into canonical AMM pools, weighted by the relative market-driven valuations of the various subnet tokens.

Dynamic TAO's design - with core token emissions bonded into protocol-owned liquidity against an ecosystem of related subnet tokens - is the closest we've seen to the end-state of DePIN tokenomics. It allows for extreme simplicity at the core protocol level, since the only job-to-be-done is to govern the process for creating/destroying subnets and weighting rewards between them, pushing all complexity out to the subnets. We expect other leading DePINs (e.g., Helium) will adopt similar structures over time.

Dynamic TAO drives defensibility by creating perfectly-competitive markets among each type of network participant. All participants must stake TAO to earn a share of emissions. Emissions follow a Bitcoin-like schedule, with 4-year halvings and a 21 million supply cap. As individual miners/validators/subnets fall behind their peers, they earn a smaller share of declining emissions and eventually fail to earn above their cost of capital relative to their staked TAO requirements.

- [Miners](#) aggregate various resources needed to produce machine intelligence (e.g. compute, data, electricity) and respond to requests from validators. Miners who lack access to low-cost inputs and/or the technical capabilities to produce useful intelligence from those inputs will quickly see rewards competed away.
- [Validators](#) query miners, rank the responses according to criteria set by subnet creators, and monetize the responses offchain. Validators that are unable to develop viable (offchain) businesses will fail to earn above their cost of capital, and eventually leave the network.
- [Subnet creators](#) design the incentive mechanisms that validators use to evaluate miner responses. Successful subnets will allow for many validators to build profitable, real-world businesses on top of the intelligence created by miners. Subnets where responses are not relevant to large end-markets, or where miners earn rewards for low-quality responses, will see their rewards competed away.

How does Dynamic TAO drive defensibility? As competition on Bittensor intensifies, it becomes impossible for any single actor to be competitive at designing, producing, *and* monetizing machine intelligence. Validators own distribution, but cannot leave Bittensor without cutting off their customers' access to the most competitive models; miners own infrastructure, but without Bittensor they cannot monetize it across hundreds of AI apps at once; subnets own the incentive mechanism, but without Bittensor they cannot get it into the hands of the best AI infrastructure providers and app developers. Therefore, we believe the success of a single subnet is a sufficient condition for the success of TAO.

3. Is a successful outcome already priced in?

Bittensor is currently the #32 highest-ranked token on [Coingecko](#) with a \$3.4B market cap and \$10B FDV. It goes back and forth with \$ASI - the token merger of FetchAI, Ocean Protocol, and SingularityNet - as the most valuable token in decentralized AI. The high valuation is primarily a function of the core community's unwillingness to sell, given that Bittensor remains a non-consensus bet across crypto funds. Public supporters include DCG, Polychain, Firstmark, Canonical, OSS, Syncracy and Collab+Currency. Public detractors include investors and researchers from [Van Eck](#), [Multicoïn](#), [Compound](#), [Galaxy Digital](#), [Delphi Digital](#), [M31](#), and [Blockworks](#). The most common criticisms are some combination of: 1) Bittensor's consensus mechanism is ultimately subjective, 2) the network does not directly incentivize "real" AI, and 3) TAO is an AI "memecoin" with no fundamentals.

In fact, Bittensor's purported weakness - lack of objective consensus - is exactly what makes it special. The ability to mobilize hundreds of millions of dollars worth of AI infrastructure and developer attention on-demand massively expands the surface area of what creative entrepreneurs can accomplish. For any problem that's sufficiently well-defined, entrepreneurs can get double-digit monthly efficiency gains. As the network attracts more resources, and learns to use those resources more efficiently, the applications built on top of Bittensor will rival and eventually beat those built by the biggest centralized AI companies.

Looking at the current subnets with momentum, we expect Bittensor subnets to power apps with 1 million end-users by the end of 2024, more than 10 million end-users by the end of 2025, and more than 100 million users by the end of 2026. Following the growth, investors' consensus view on Bittensor will shift from *AI memecoin with no onchain consensus* to *the only smart contract platform capable of incentivizing intersubjective value*; new cohorts of app developers will be inspired to create products "Only Possible on Bittensor"; and the fair value of the Bittensor network will re-rate to levels comparable to leading smart contract platforms and AI companies. The Bittensor network gets stronger as it gets bigger.

Over a longer time frame, Bittensor has the potential to be bigger than anything in crypto today—potentially even bigger than Bitcoin. While the latter was borne out of over-regulation of money, the former is borne out of over-regulation of intelligence. At scale, Bittensor has the potential to become the most liquid AI asset in the world, the highest-paying "employer" of AI talent in the world, the largest computing cluster in the world, and - eventually - the most powerful intelligence in the world.

There are many criticisms of Bittensor in its current state which are accurate. There's a meaningful [centralization](#) at the physical layer which will require transitioning to proof-of-stake (expected 2025). Several current or past subnets - especially those run by venture-backed teams - existed solely for the purpose of extracting value from emissions (e.g., Nous Research or MyShell). There are weaknesses in the consensus mechanism which make the network far less efficient than it could be (e.g., [weight-copying](#) by validators). There's been several [chain halts](#) for bugs and/or network attacks in the past year. There's fairly limited institutional support for custodying and/or staking TAO. However, we don't see any of these challenges - or their sum - as insurmountable for an ecosystem with as strong a flywheel as Bittensor.



*** DISCLAIMERS ***

This letter has been prepared by EV3 Ventures LLC ("EV3 Ventures"). EV3 Ventures is providing this letter for informational and discussion purposes only. The views and opinions expressed in this letter reflect the views and opinions of EV3 Ventures as of the date hereof. EV3 Ventures reserves the right to change or modify any of such views or opinions at any time and for any reason and expressly disclaims any obligation to correct, update, or revise the information contained herein or to otherwise provide any additional materials to any recipient of this letter.

All of the information contained herein is based on EV3 Ventures' independent research and analysis and publicly available information. EV3 Ventures does not make any representations regarding the accuracy, completeness, or timeliness of any third party statements or information contained in this letter. This letter does not purport to contain all information that may be relevant to an evaluation of any investment in the digital assets described herein, any securities issued by any issuer described herein, or any other investment opportunity. The contents of this letter are not intended to be and should not be viewed as or relied upon as financial or investment advice.

Nothing in this letter should be taken as an indication of any current or future trading or voting intentions, which may change at any time and from time to time, and which are reviewed on a continuing basis by EV3 Ventures and determined based on various factors EV3 Ventures considers relevant. EV3 Ventures expressly disclaims any obligation to notify or update any person of any changes to such trading or voting intentions.

Certain statements contained in this letter are forward-looking statements including, but not limited to, statements that are predictions of or indicate future events, trends, plans or objectives. Undue reliance should not be placed on such statements because, by their nature, they are subject to known and unknown risks, assumptions, and uncertainties. Specific forward-looking statements can be identified by the fact that they do not relate strictly to historical or current facts and include, without limitation, words such as "may," "will," "expects," "believes," "anticipates," "plans," "estimates," "projects," "targets," "forecasts," "seeks," "could," "should" or the negative of such terms or comparable terminology. There can be no assurance that any idea or assumption herein is, or will be proven, correct. If one or more of the risks or uncertainties materialize, or if EV3 Ventures' underlying assumptions prove to be incorrect, the actual results may vary materially from outcomes indicated by these statements. Accordingly, forward-looking statements should not be regarded as a representation by EV3 Ventures that the future plans, estimates, or expectations contemplated will ever be achieved.

This letter does not constitute an offer to sell, a solicitation of an offer to buy, or a recommendation to buy, sell, or otherwise transact in any security, including, without limitation, any interest in any pooled investment vehicle, fund, or other account managed by EV3 Ventures (an "EV3 Fund"). Any offer to purchase an interest in an EV3 Fund may only be made at the time a qualified offeree receives the confidential memorandum and other offering documents. No agreement, commitment, understanding, or other legal relationship exists or may be deemed to exist between EV3 Ventures and any other person solely by virtue of such person's direct or indirect receipt of this letter.