
Compute for the
frontiers of artificial
intelligence

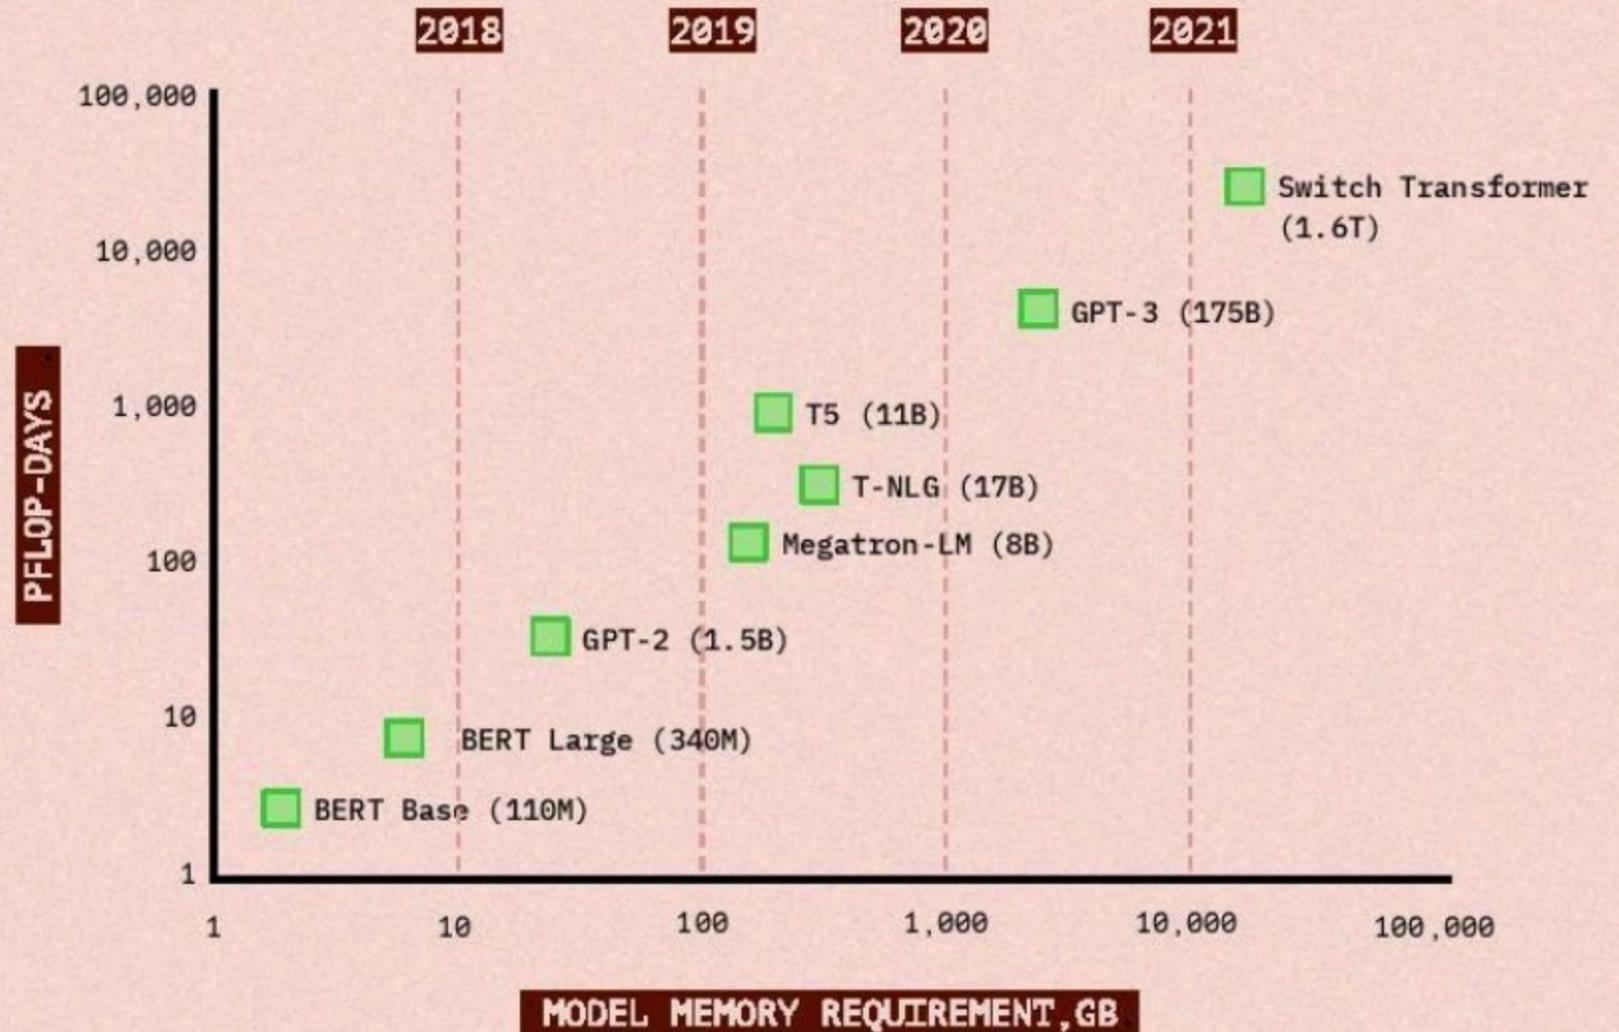


Demand for AI compute power is **exploding**

The computational complexity of state of the art AI systems is doubling every 3 months

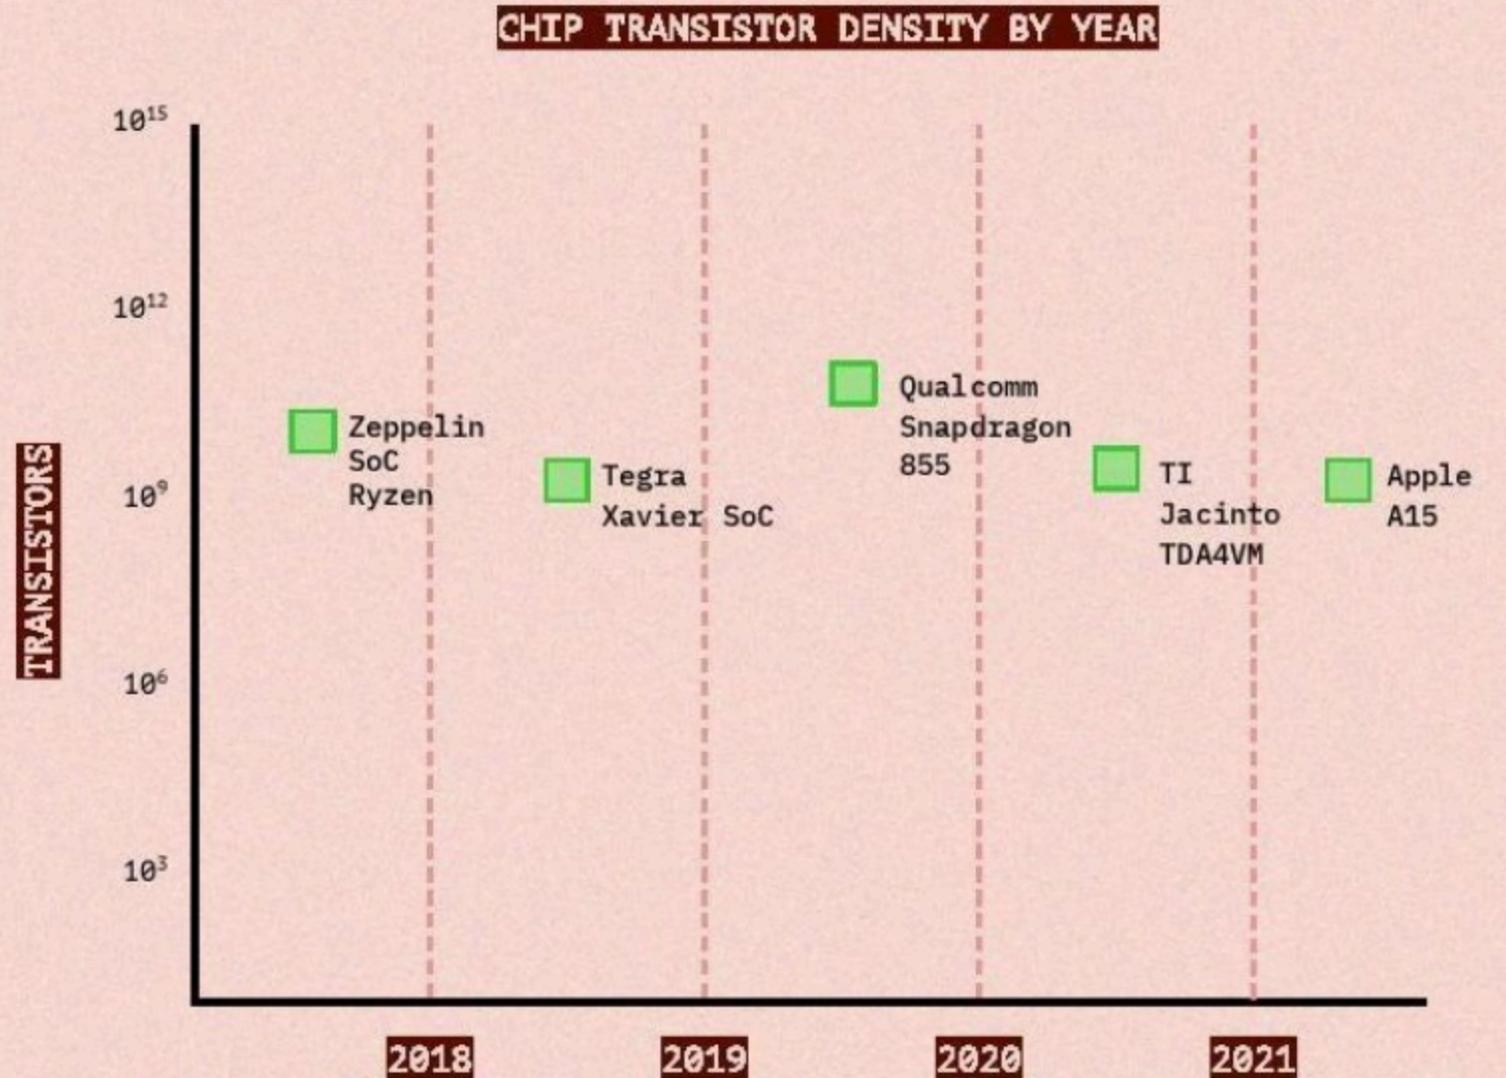
A \$181bn opportunity by 2031

STATE-OF-THE-ART MODEL TRAINING TIME BY YEAR



Supply is **Lagging**;
users must pay up
or scale down

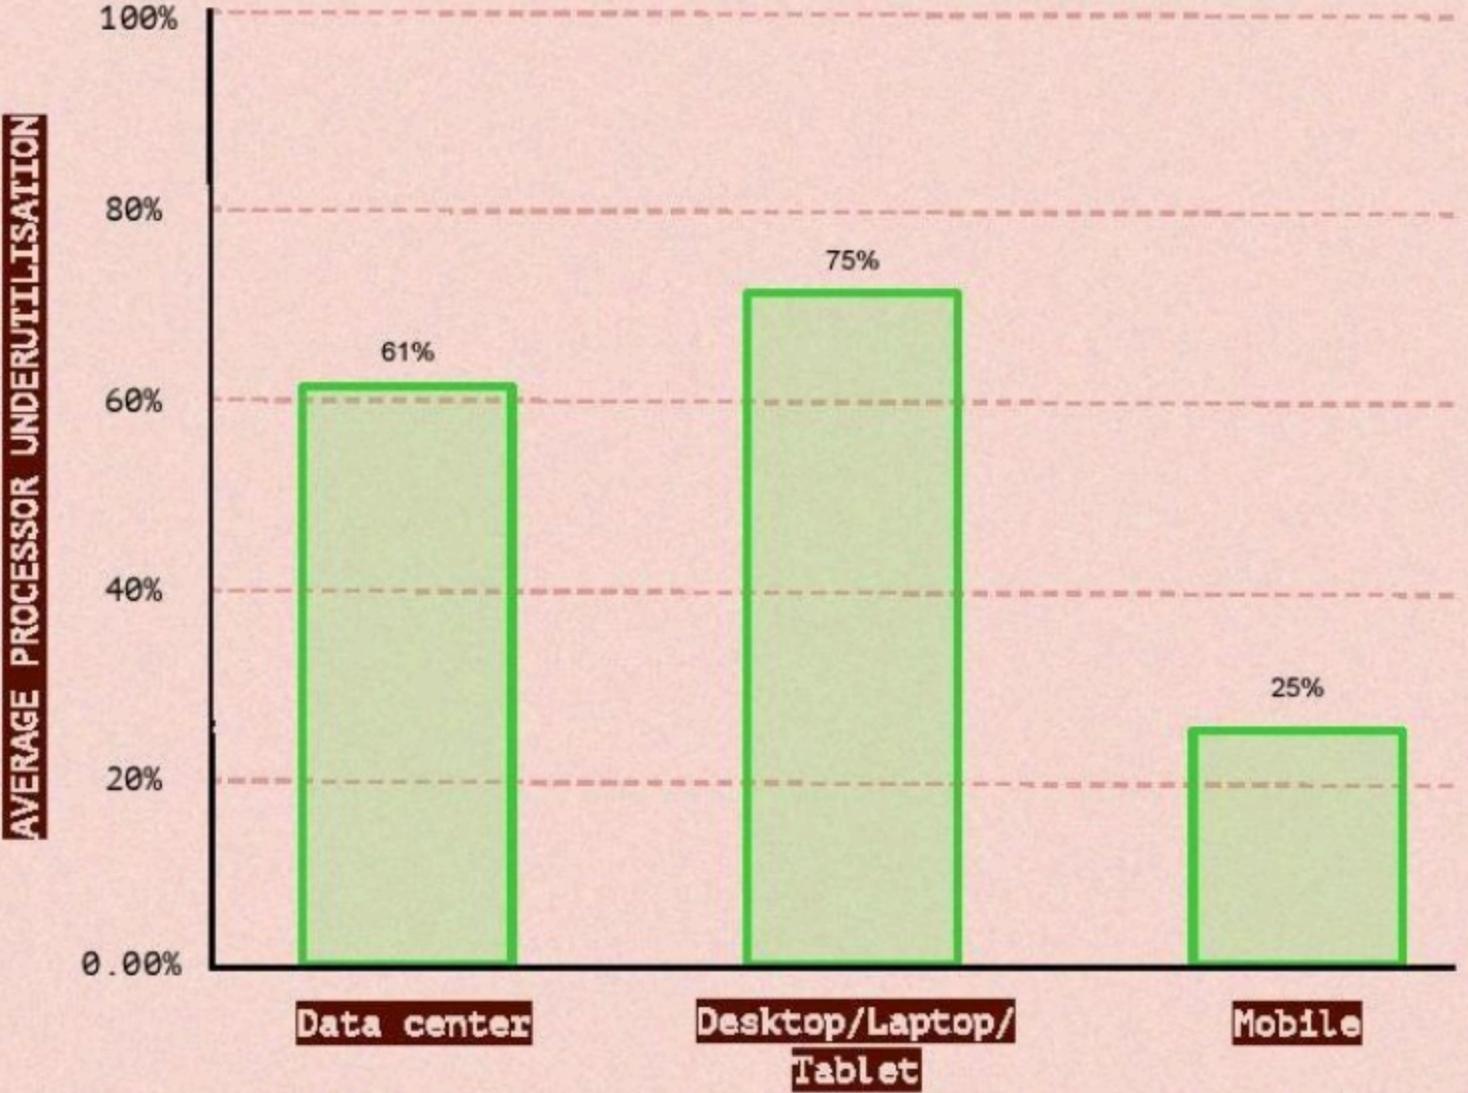
- Over 60% of general compute is wasted
- Asymptotic [chip](#) performance gains
- Ballooning demand from [new sectors](#)
- Geopolitics leading to [stockpiling](#) of chips
- Centralised compute options are either:
 - ◆ Expensive oligopolists (e.g. [AWS](#)); or
 - ◆ Limited in supply due to [incentives](#)



Over **60%** of general compute supply is wasted

- Data center hardware is underused
- Personal computers sit idle
- Mobile phones have huge potential for overnight utilisation
- ETH1 miners will soon be searching for yield

Problem: all of this compute is untrusted



Gensyn is the L1 **trustless** protocol for machine learning computation

Read more in our [Litepaper](#)

Fundamental feature

- Work verification: probabilistic method for approximate verifiability and deterministic, game-theoretic, challenge mechanisms for computational work verification

Key features

- Market: tokenized market for compute with up-front cost estimation, foundation models for warm starts, and on-demand execution
- Ex-ante work estimation: neural network graph unravelling using intermediate representation with specified loop parameters to prevent the [halting problem](#)
- Privacy: functional encryption layer allows computation over private datasets
- Parallel optimisation: implementation of state-of-the-art distributed optimisation research for parallel training

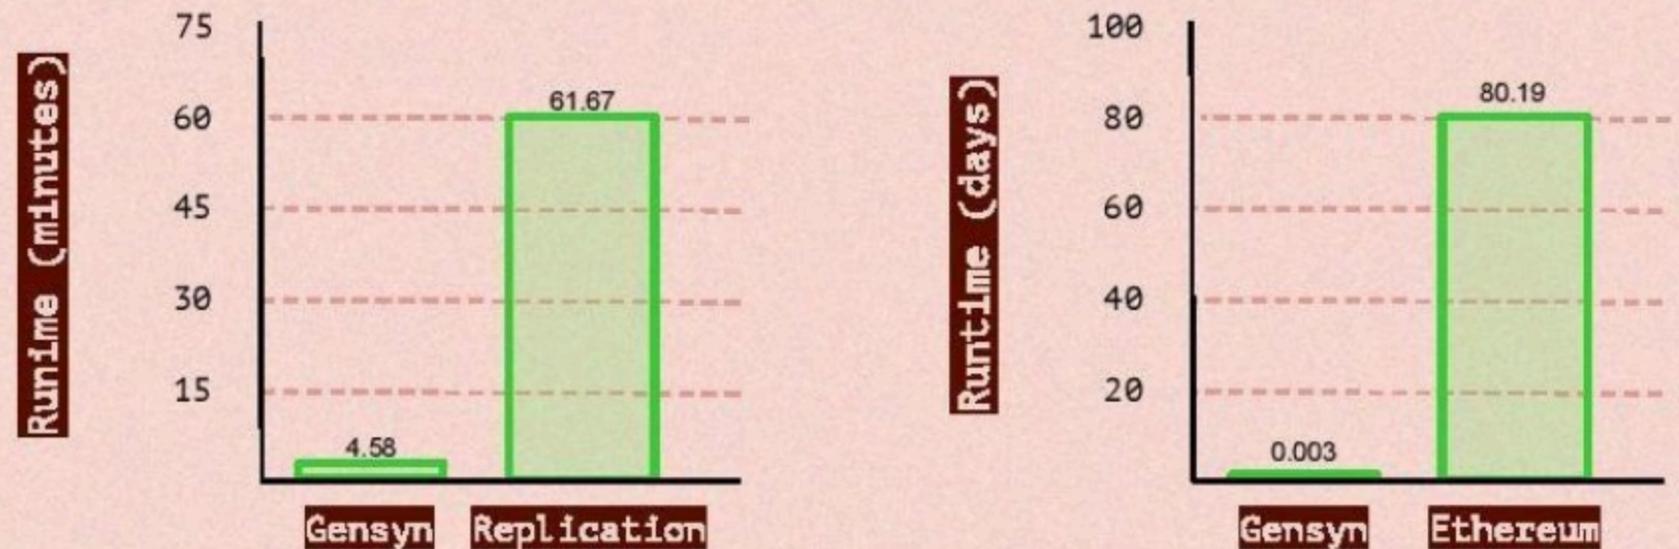


A **first of its kind** scalable deep learning verification mechanism

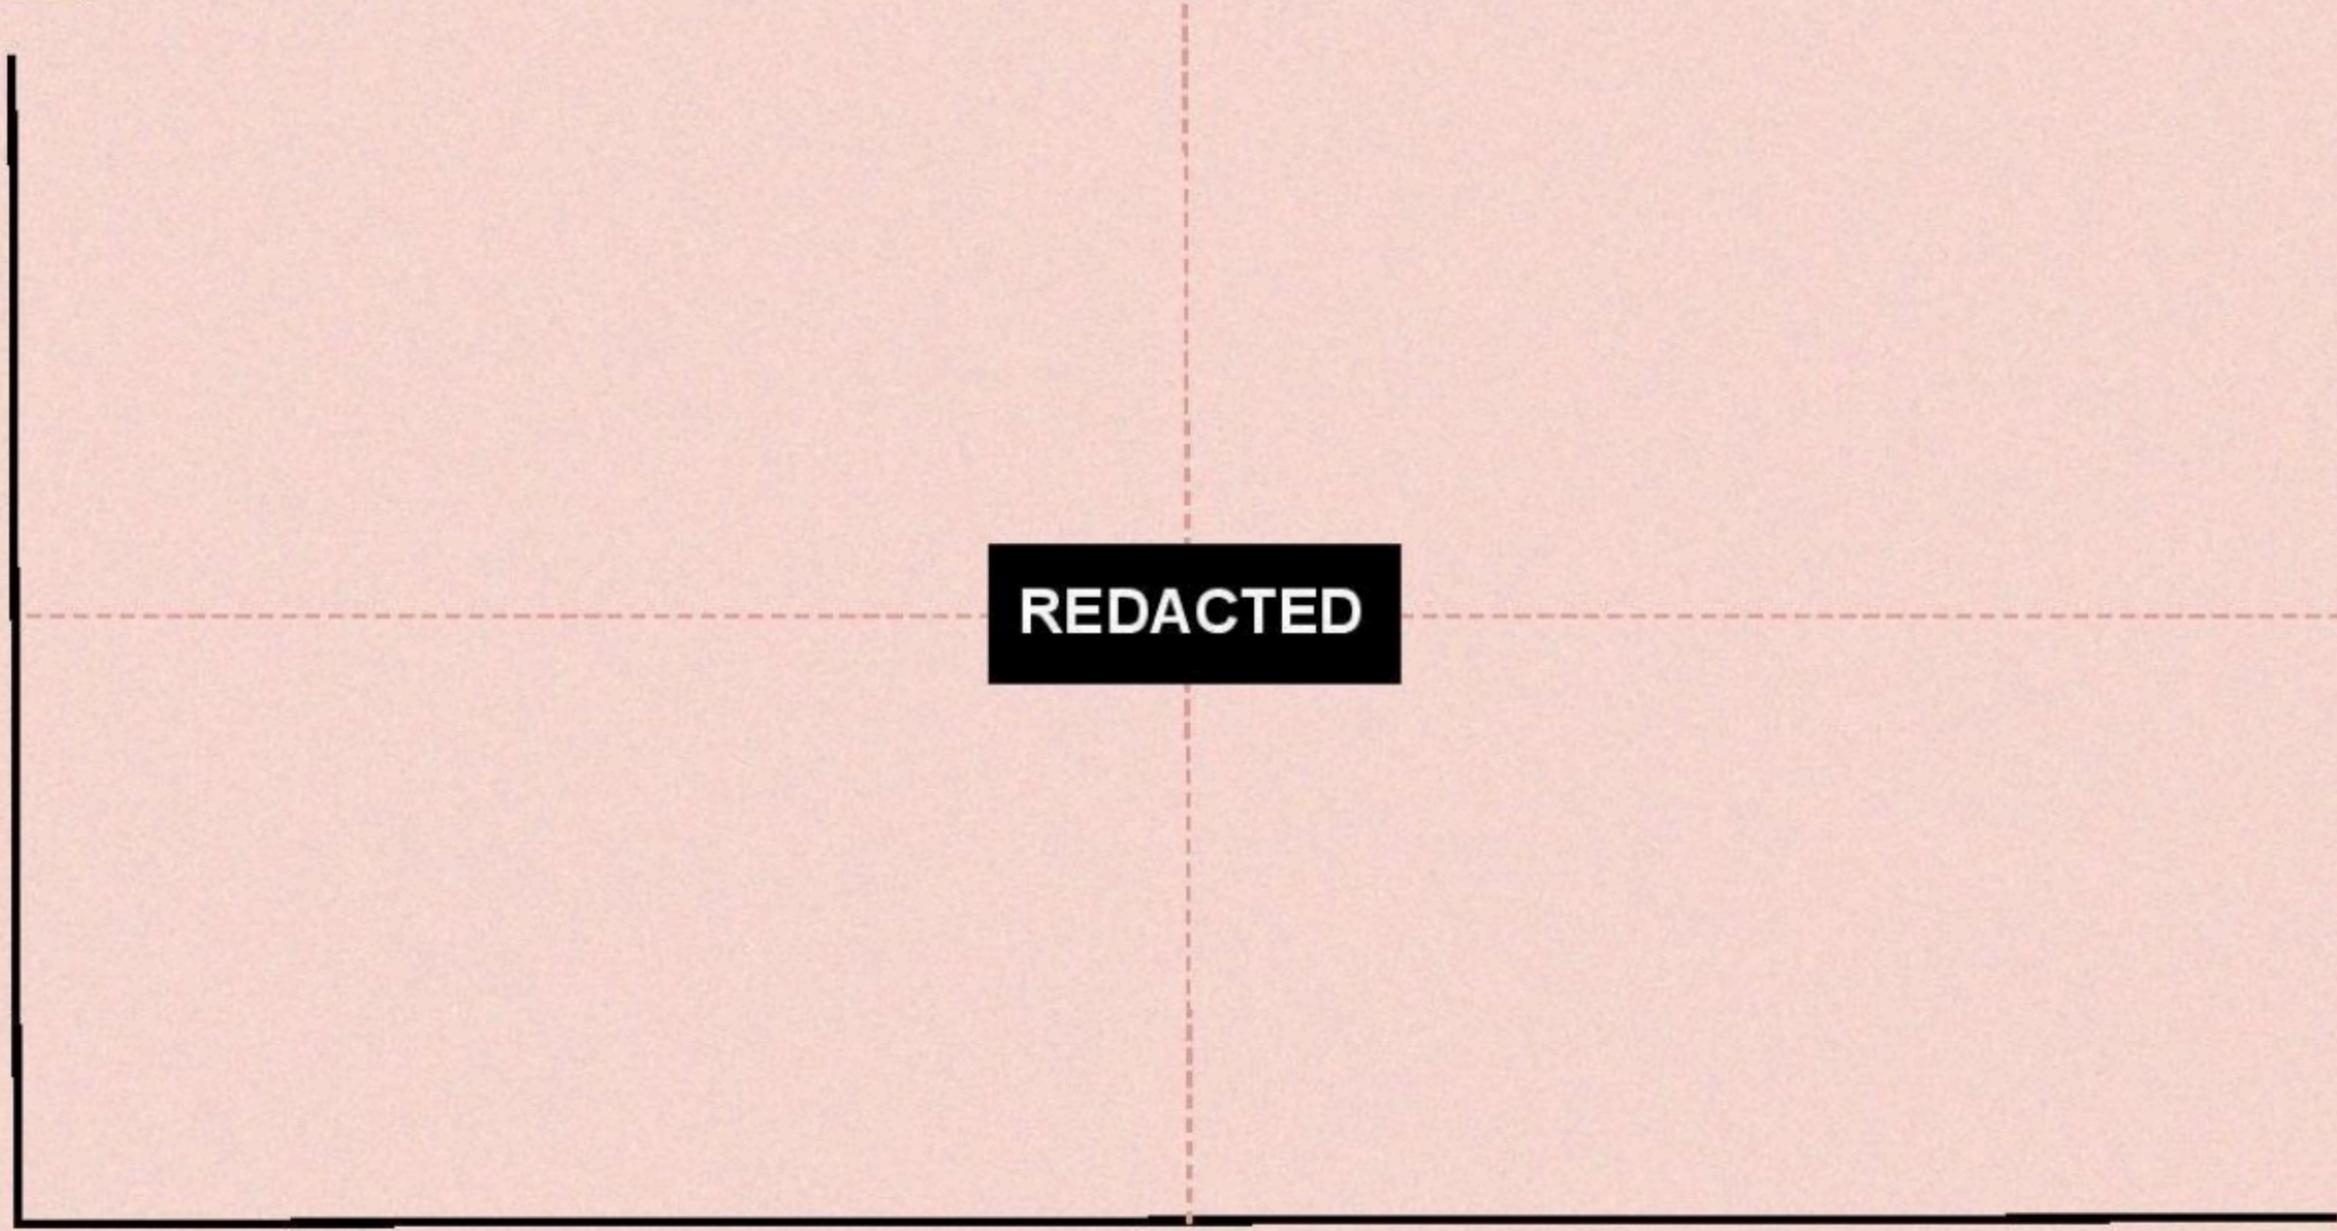
1,350% faster vs the best competition

- **Probabilistic method for approximate verifiability:** using the metadata from gradient based optimisation methods to check that a model has been trained as requested
- **Deterministic challenge mechanisms:** a graph based pinpoint protocol for contract arbitration between trainers and verifiers
- **Game Theory:** a Truebit-style incentive game that produces honest participation from financially rational actors through slashable deposits and reward bounties

MNIST image classification model verification time



**AI COMPUTE
SCALE**



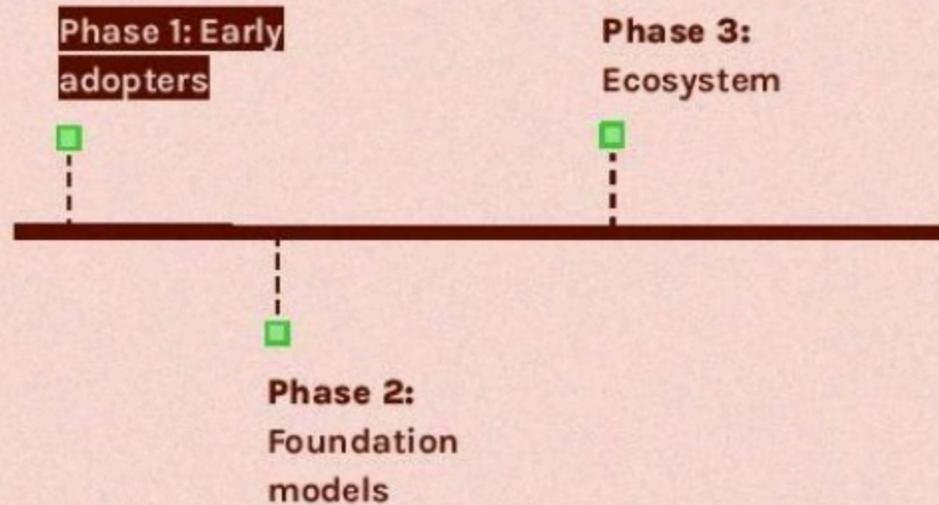
REDACTED

COST EFFICIENCY

* approximate cost of maxing out deep learning training on Nvidia V100 (or equivalent) per hour



Designed for **ML researchers and engineers** in academia and industry



Phase 1: Early adopters

- User research: we've spoken with >150 ML researchers and engineers from academia and Seed to Series B companies about their requirements, such as:

REDACTED

- Community: from those conversations, we've built a curated Discord community of 50 early adopters who are keen to train new models on the network immediately

Phase 2: Foundation models

- End users: ML researchers and engineers fine-tuning huge, pre-trained models (BERT, GPT-3, DALL-E) on new datasets

Phase 3: Ecosystem

- End users: SDEs using applied AutoML Dapps on top of the network without requiring in-depth ML or protocol knowledge



Decentralising
the market for
compute power
generates a
\$181bn market
cap by 2031

Read more in our [Token Valuation Model](#)

- **Artificial General Intelligence (AGI) advancement:** Dramatically advances ascent to AGI by fuelling the design and training of collaborative [foundation models](#)
- **Cost effective:** both high in scale/supply given access to the global compute stack and cost effective as the price of compute is moved closer to cost-price of processor operations
- **Environmental:** [greener](#) return on processor work by attracting GPU (or custom processor) power away from PoW scheme; as well as reaching consensus via PoS
- **dApp enablement:** facilitates higher layer functionality (e.g. a decentralised [DataRobot](#) or [BenevolentAI](#))
- **Censorship resistant:** a protocol governed by elected council with no proof of identity

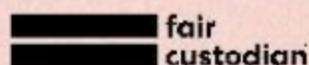


World class team
and investors
with deep AI
experience

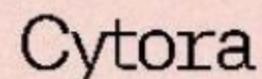


Founders:

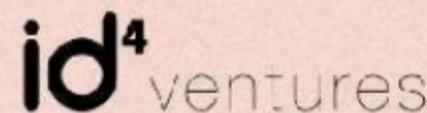
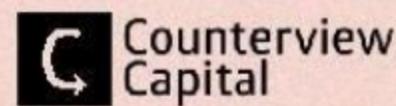
Ben: PhD in ML
optimisation, ex founder of
anonymous digital identity
startup



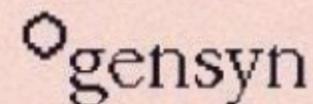
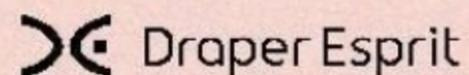
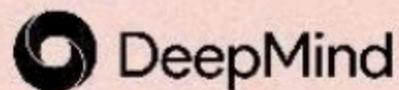
Harry: ex Head of Data Research
at ML risk-pricing startup,
background in applied
econometrics



Pre-seed investors:



Angels from:



Contact :

founders@gensyn.ai





Browse the best pitch deck examples.

Brought to you by bestpitchdeck.com — the world's largest library of pitch decks: hundreds of winning presentations from leading startups, updated every week.

[Read more →](#)

Follow us [@pitchdecks](#)    

