

∴ nineteenai

Decentralised AI Inference at Scale

namoray bittensor



The **Problem** with Centralised AI

Ⓒ Single point of **failure**

Centralised control can lead to a single point of failure, making systems more vulnerable to errors or attacks.

⊘ Lack of **transparency**

Companies like OpenAI and Anthropic often degrade the performance of their existing models after a short period without notifying users, leaving you unaware of what's actually being ran.

🔒 **Privacy** & Data Protection

Closed source models put your personal data at risk, allowing single entities to use or expose your private data at their whim

🌐 **Cost** at Scale

Leading AI companies charge high prices to maximise profits, rather than setting prices based on real market demand, making it unaffordable for many

Subnet 19: The **Solution**

∴ nineteenai bittensor

Decentralised System

Subnet 19 is the definition of decentralised AI inference, leveraging a network of miners to run Generative AI models, meaning that there is, and will never be a centralised server to attack.

Open Source, Transparent Code

Subnet 19 is all built on a transparent, open-source code base that promotes a more inclusive and accessible ecosystem.

Privacy & Security Baked In

We're not just giving access to cutting-edge AI models, we're giving back control. Every validator provides their own access and requests are distributed across different unique entities, meaning that there are no centralised databases of responses saved for training purposes












The **Lowest** Cost at Scale

The price of our inference is significantly cheaper than comparable centralised solutions, with current API providers built on-top of Nineteen offering over 60x cheaper prices than similar centralised offerings. This changes the game for the types of products that can be built with ai by both startups and established firms.



How do we compare?

Subnet 19 is at the epicenter of the AI revolution, where the latest advancements in AI inference are being made. As the industry hurtles towards a future that's more automated, more personalized, and more efficient, we're proud to be leading the way.

	together.ai	 Fireworks AI	 OctoAI	 nineteen.ai
Decentralisation	✗	✗	✗	✓ Entirely Decentralised
Speed	Fast  Standard Speed	Fast  Standard Speed	Fast  Standard Speed	Faster  Decentralised network, with the highest potential speed.
Types of Models	Text, Image	Text, Image	Text, Image, Video	Text, Image <small>More Coming Soon</small>
Scalability	 Global	 Global	 Global	 Global



Open Source Large Language Models have the edge.

Subnet 19 miners provide access to the newest SOTA open sourced models like Llama3.1 8B & Llama3.1 70B.


One of the many reasons to use subnet 19s LLMs, TTFT ~3x quicker than closed sourced models


Time to first token GPT-4o

2.0 Seconds

Time to first token Llama 3.1 8b

0.7 Seconds





All in one subnet.

Subnet 19 provides access to a full set of the best LLM and image models.

Large Language Models

Image Models


Other Models

Coming Soon

The screenshot displays the AI interface with the following components:

- LLM Interface (Left):** Shows 'Llama 3.1 70b' with a 'Meta' tag. It features a performance graph and a table with the following data:

Tokens/s (Average):	80	Tokens/s (Max):	120
Knowledge Cutoff:	2023/12	Inference:	Subnet 19
- Image Generation Interface (Right):** Shows 'Flux' with a '8.16 Generation Time (s)' indicator. It displays two generated images of a penguin in a top hat holding an umbrella in front of Big Ben. Below the images is a text input field containing the prompt: 'a penguin in a top hat and holding a umbrella, standing in front of big ben in'.



World-class text models

Subnet 19 enables decentralized inference for the world's leading open-source large language models.



Llama 3.1 70b

94 Tokens/s

✓ 100% Proof of Model Authenticity

0 Degradation in quality



Llama 3.1 8b

188 Tokens/s

✓ 100% Proof of Model Authenticity

0 Degradation in quality

World-class image models

Subnet 19 provides decentralized inference for some of the world's best image models available.

Flux-Schnell

The master of text

Proteus

5/5 quality



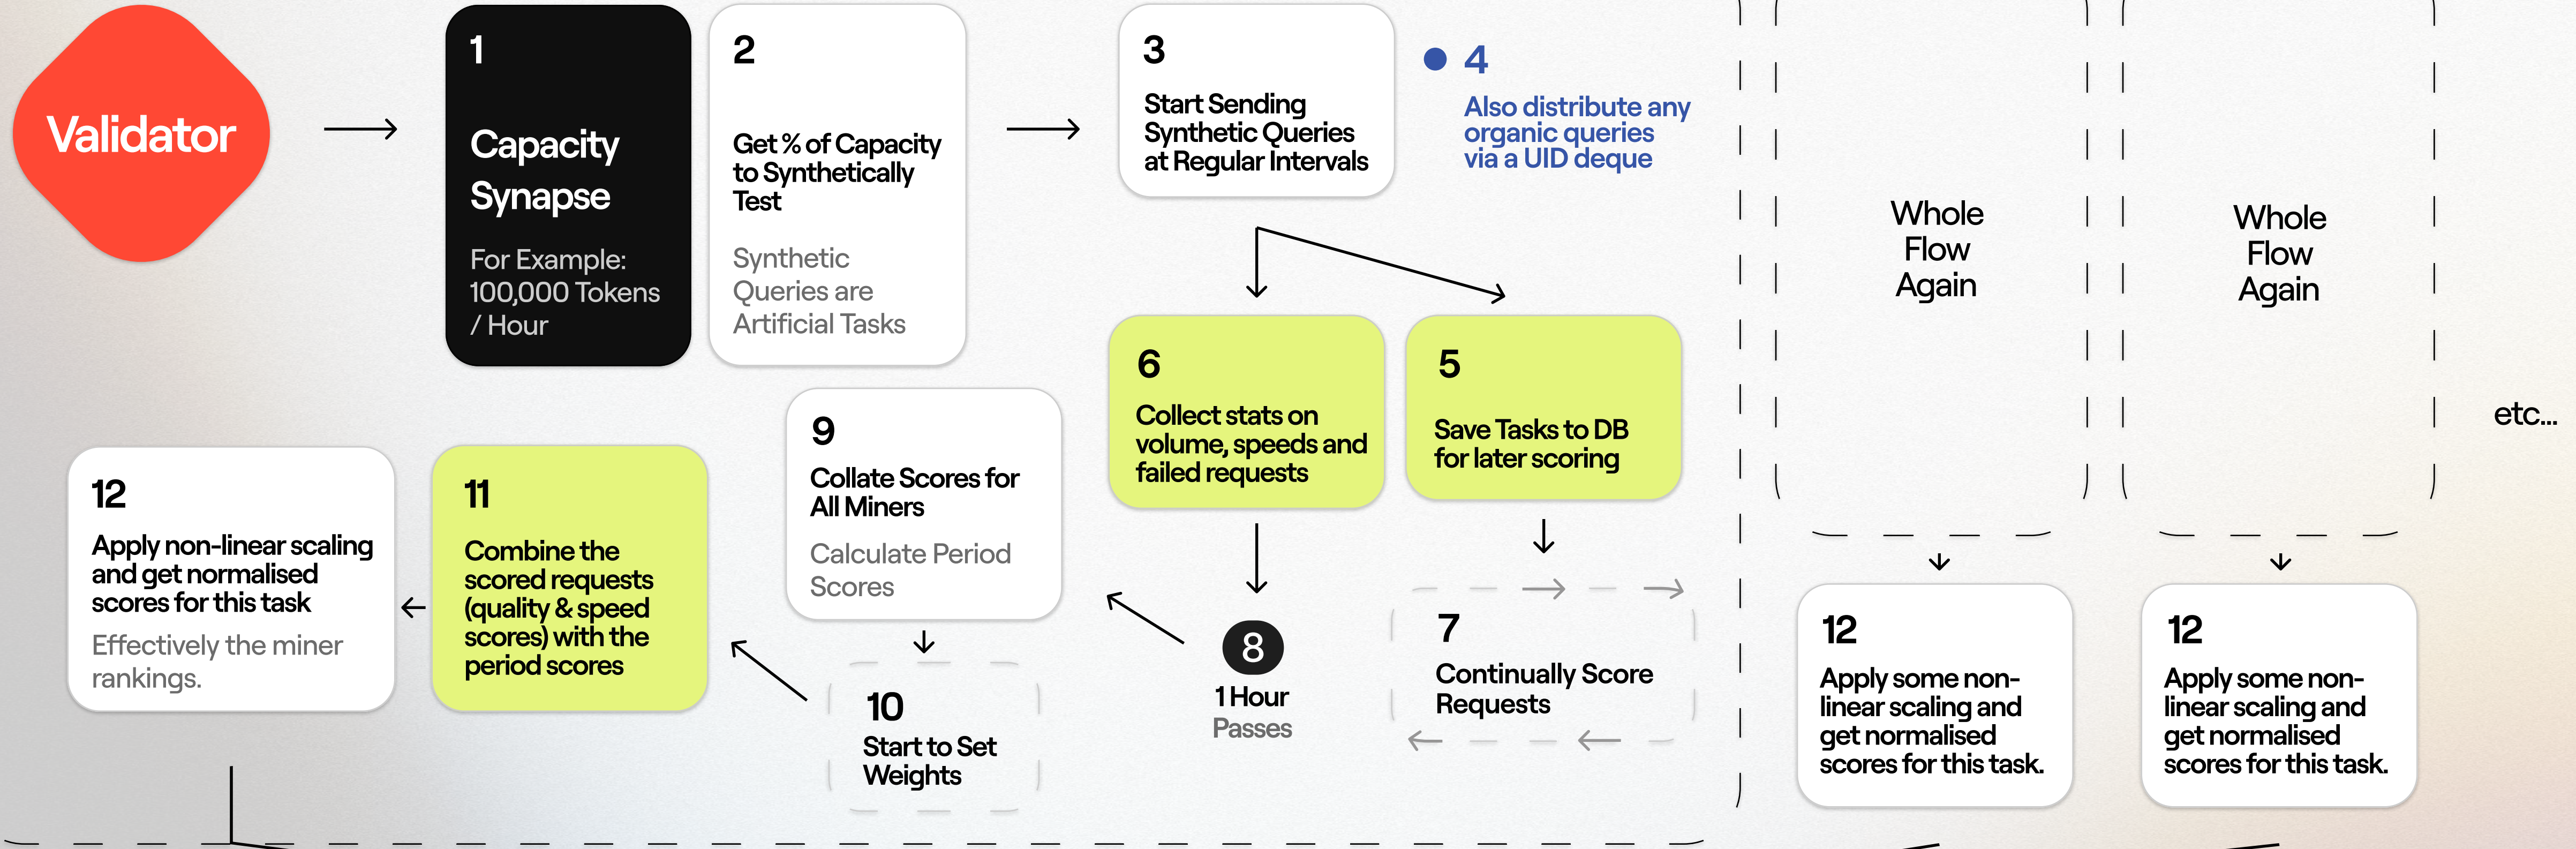
Dreamshaper

You've never seen realism like this



⋮ Mechanism


 Chat Llama 3.1 70b



13 Combine all scores with the task weights to get a final incentive (13)


14 Set Weights with Retries

Why NineteenAI on Bittensor?



Completely Decentralised

100% of everything in our subnet is open-sourced and decentralised.




Most Valuable Digital Commodity

Revenue generation and real-world use are not an afterthought for us. We have the most traffic out of any bittensor-built product.




120+ Tokens/s for Top-performing Models

Speed is crucial in the world of inference, and NineteenAI is no different. On top-performing models like Llama 3.1 70b we achieve over 120+ tokens/s.



Community Driven

Proposals for every update followed by in depth discussion ensure constant improvements.



Fastest Iteration Time

We are always improving, and never stagnant. We ship quickly but iterate even quicker, and development is very active.



The Highest-quality Bittensor products are built on NineteenAI

Millions of messages and hundreds of thousands of users.

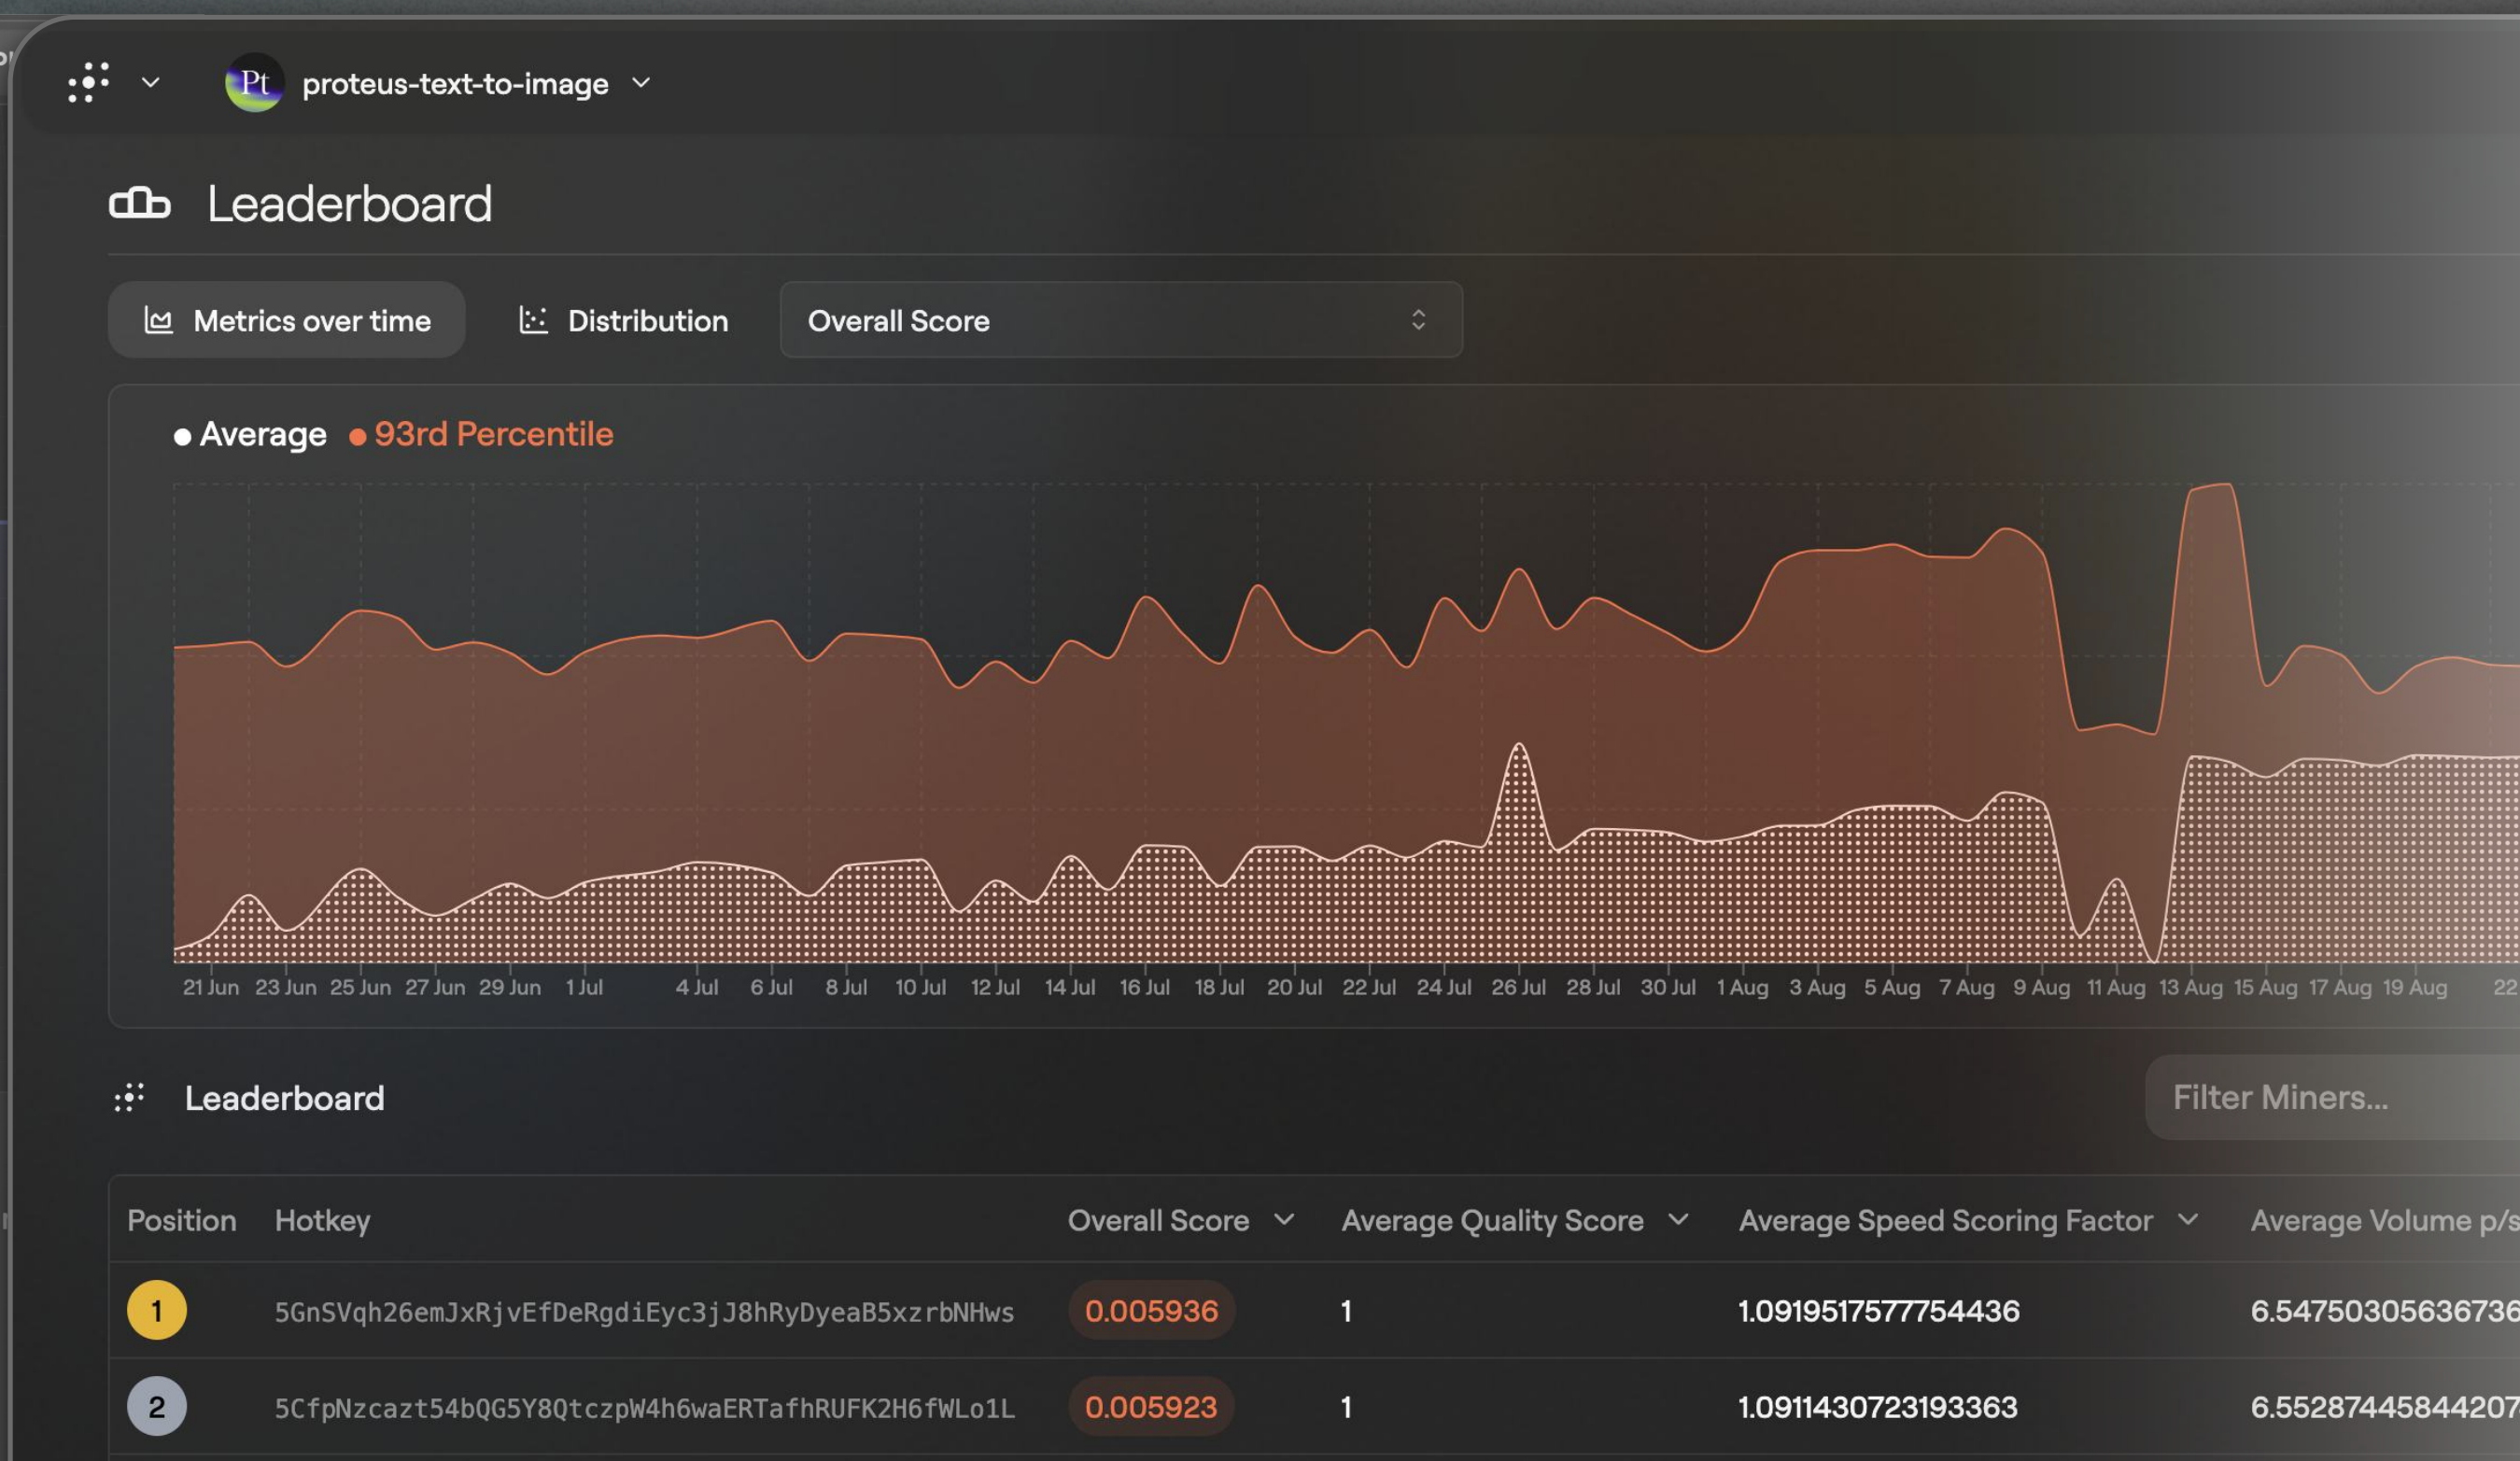
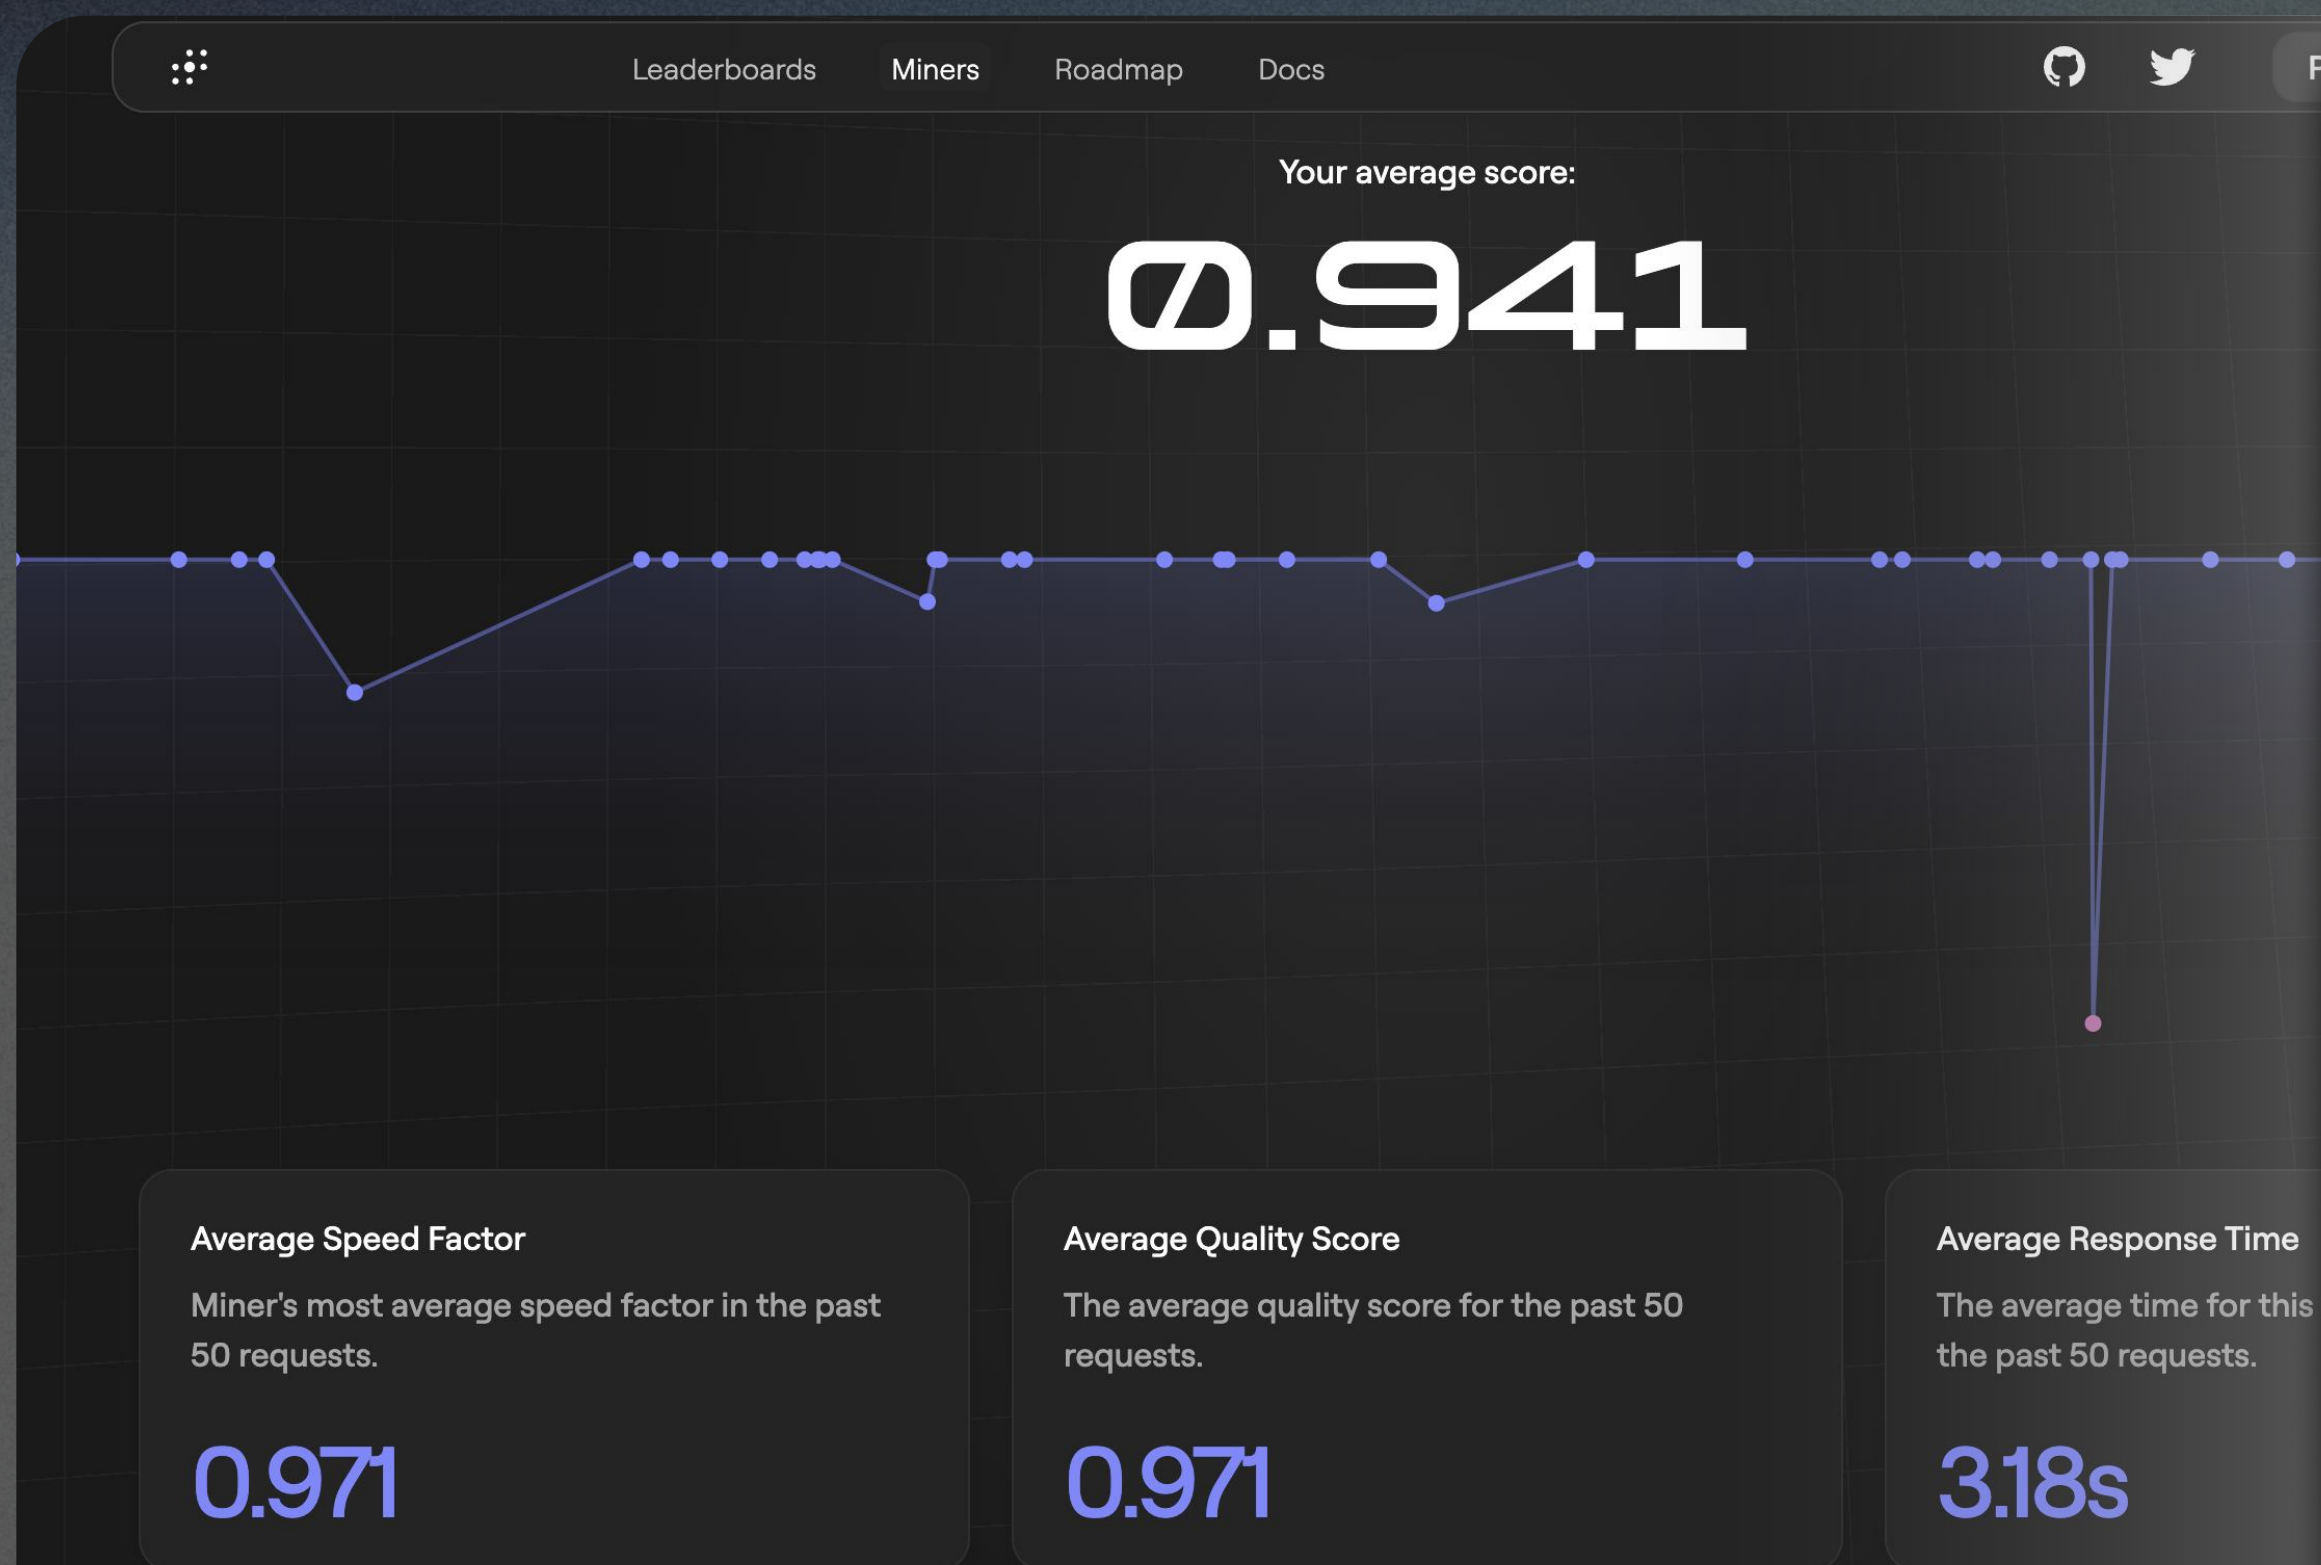


For Miners

- Task weightings
- Miner speed and volume scores

For Validators

- Subnet metrics over time
- dAPI (Decentralised API Access)





Build *Lightspeed* Bittensor Subnets with Fiber

Nineteen has rebuilt the Bittensor package from the ground up with Fiber.


Measured in hundreds of lines of code instead of hundreds of thousands, Fiber focuses on what truly matters for subnets,

offering enhanced security, DDOS resistance, and portability to all your favorite languages - straight out of the box.

Bittensor-less* subnets on Bittensor are here.

namoray bittensor

*A subnet built with Fiber no longer requires the Bittensor Python package to function.





Our Journey + Roadmap

20th Dec 2023

Vision 0.1.0

Giving Eyes to Bittensor. The world of images, unlocked. Recognition, Embeddings, Object detection, Segmentation

24th Dec 2023

Vision 1.0

Text to Image, Image to Image, Upscaler and Inpainting are now available

19th February 2024

Vision 2.0

Validators can sell their organic traffic in a completely decentralised way. All decentralised, no API keys. We have SDXL Turbo, Kandinsky, Inpainting, Scribble, Controlnets, Upscaling, Image-To-Image, Uncrop support, the list goes on!

11th Mar 2024

OpenVision

Sn19s own SOTA image model, trained on midjourney style images

5th Sep 2024

Fiber

Introducing Fiber - our fully reworked Synapse, Dendrite and Axon mechanism. Simplifying and speeding up the development of the whole Bittensor ecosystem

28th Aug 2024

NineteenUI

At Nineteen, we now have our own front-end UIs for all our services. Completely free to use, without signup or registrations.

10th Jun - 9th Aug 2024

Vision 4.0

Full decentralisation, working our way to 1uid per entity & new SOTA's in town: Llama 3.1 70B, Llama 3.1 8B, and Flux Dev

4th Apr 2024

Vision 3.0

Open source LLMs arrived on SN19, starting with Mixtral 8x7b

More types of models

Coming Soon

Coming Soon

To NineteenAI

Coming Soon

To NineteenAI

Coming Soon

To NineteenAI



nineteenai

namoray bitensor