Zeus: Decentralized Climate Forecasting

Whitepaper v0.1 – February 7, 2025

Ørpheus AI

Authors

Travis van Cornewal   Egill de Visser

Abstract— **Accurate climate forecasting is essential for mitigating risks and optimizing operations across sectors such as agriculture, transportation, and energy management. Traditional numerical weather prediction systems rely on complex modeling and high computational resources, limiting their adaptability and scalability. Machine learning-based weather prediction offers a promising alternative by leveraging historical data to improve forecast accuracy and efficiency. Zeus introduces a decentralized climate forecasting framework built on Bittensor's blockchain, fostering a competitive ecosystem where machine learning engineers fine-tune state-of-the-art models under defined constraints. Zeus provides a modular and scalable approach that can integrate environmental variables over time. A novel incentive mechanism ensures fair evaluations through difficulty-adjusted rewards, while out-of-sample forecasting prevents overfitting and enhances predictive reliability. By decentralizing climate forecasting, Zeus democratizes access to advanced modeling techniques, ensuring continuous innovation, transparency, and resilience against climate change.**

1 Introduction

Climate forecasting plays a critical role in a wide array of sectors, from transportation to public safety, agriculture, and energy management. Accurate climate predictions enable governments, businesses, and individuals to make informed decisions that minimize risks and optimize operations. As climate change continues to alter weather patterns, the need for accurate, timely, and robust climate forecasting systems is becoming even more critical, highlighting the importance of advancing forecasting techniques and technologies to better address these emerging challenges.

The conventional end-to-end weather prediction pipeline, based on numerical weather prediction (NWP) systems, is a complex, multi-stage process that integrates diverse observational data and advanced modeling techniques. It begins with data acquisition, which involves the collection of observations from remote sensing instruments, in-situ platforms, radar systems, and radiosondes, alongside derived products such as atmospheric motion vectors. This is followed by atmospheric state estimation, where raw and processed data are assimilated with prior forecasts to create a comprehensive approximation of the current atmospheric state. Next, the forecasting stage uses this state as an initial condition to predict future atmospheric conditions through models based on fluid mechanics and thermodynamics. Finally, post-processing and downstream applications refine these predictions using statistical techniques and higher-resolution regional models, generating actionable local forecasts (Vaughan et al., 2024).

This intricate workflow demands significant computational resources and expertise, with improvements to conventional NWP methods often relying on highly trained experts developing better models, algorithms, and approximations—a process that can be time-consuming and costly (Vaughan et al., 2024; Lam et al., 2023; Price et al., 2024).

Machine learning-based weather prediction (MLWP) is rapidly advancing and presents a competitive alternative to traditional NWP systems. By training models on historical data, MLWP can capture complex atmospheric patterns that are difficult to represent with traditional equations. This approach holds the potential to improve forecast accuracy while offering greater efficiency, leveraging modern deep learning hardware instead of supercomputers to achieve favorable speed–accuracy trade-offs (Lam et al., 2023). While NWP systems like ECMWF's high-resolution forecast (HRES) have been the standard for medium-range forecasting, recent developments in MLWP have demonstrated superior performance in certain areas, including forecast skill and efficiency. These advancements have allowed MLWP not only to match, but in some cases to surpass the accuracy of traditional NWP systems (Lam et al., 2023; Price et al., 2024).

1.1 The Role of Zeus

Zeus introduces a novel approach to MLWP by leveraging Bittensor's network to create a competitive ecosystem of forecasting models. Zeus is a decentralized climate forecasting framework that initially focuses on predicting a single variable, the temperature two meters above the Earth's surface. Built with modularity at its core, the framework is designed to evolve into a more comprehensive forecasting framework. It enables the seamless integration of additional environmental variables over time. Zeus leverages a global network of engineers to distribute forecast challenges, optimizing state-of-the-art AI models for greater accuracy and efficiency.

Machine learning engineers within the Zeus subnet independently fine-tune models under defined time and hardware constraints, creating an environment where innovation thrives. By competing to refine the models, engineers adopt cutting-edge techniques that offer a competitive edge, resulting in diverse approaches to solving forecasting challenges. This decentralized architecture not only drives innovation but also ensures continuous adaptability, allowing Zeus to respond to local conditions and incorporate new data streams as they emerge.

Climate forecasting on Zeus is transparent and incentivized through Bittensor's blockchain, Subtensor. By securely recording key activities related to model development and reward distribution on-chain, Subtensor provides a transparent, immutable system of record, ensuring the integrity and public verification of Zeus's operations. This decentralized framework drives innovation through competition among engineers, who are all incentivized to contribute to the continuous improvement of climate forecasting.

2 Bittensor

Bittensor is a network that leverages blockchain technology to facilitate decentralized artificial intelligence. In this ecosystem, a subnet serves as an incentive-driven marketplace which defines the work that the participants of a subnet — miners and validators — must perform. Miners contribute by performing useful work defined by the subnet's incentive mechanism, while validators independently evaluate the task performed by the subnet miners based on predefined standards. All the participants, including the subnet owners, are rewarded with TAO, the native currency of the Bittensor ecosystem.

3 Architecture

To facilitate the decentralized climate forecasts, Zeus leverages the core principles of Bittensor network by building upon its subnet architecture accordingly. Three core roles form the foundation of the network, each playing a critical part in operational success:

- **Miners** are tasked with running forecasting algorithms that predict environmental variables at specific locations and timestamps.
- **Validators** challenge miners with subsets of environmental data and evaluate their performance.
- **Subnet owners** are responsible for managing the operational settings and governance of the network. They ensure that the incentive mechanism promotes alignment with network objectives while also fine-tuning system parameters to enhance performance. Their role is to keep the network evolving and prevent gamification.

3.1 *Workflow*

The Zeus subnet employs a workflow that can be broken down into three stages:

- **Stage 1:** Validators sent input data to miners.
- **Stage 2:** Miners respond with a forecast.
- **Stage 3:** Validators evaluate and reward the forecast.

3.2 *Forecast challenge*

Miners are challenged by validators to forecast environmental variables for a specified region, requiring hourly forecasts for a set of latitude-longitude coordinates over a given time window. Miners receive a 2D geographical grid, a start time point and a forecast window as input data. Then, validators require miners to return a forecast for each unique set of time and location. Initially, the focus is on forecasting the *2m temperature* variable in degrees Kelvin, denoted as T_{2m}, representing the temperature two meters above the Earth's surface.

3.2.1 *Geographical grid A*

Miners are given a matrix representing a geographical grid containing latitude and longitude coordinates. Let A be an $m \times n$ matrix denoted by:

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \dots & a_{mn} \end{bmatrix}$$

Here, both m and n are sampled independently from a discrete uniform distribution

$$m, n \sim U\{a, b\}$$

with lower bound $a = 4$ and upper bound $b = 12$, where each unit step corresponds to 0.25 degrees in latitude or longitude. The geographical location of matrix A is selected randomly where the elements a_{mn} are evenly spaced at 0.25 degrees intervals, preserving the structured nature of the geospatial representation. Each element a_{mn} in A corresponds to a unique set of coordinates (i, j) where i represents the latitude coordinate and j represents the longitude coordinate, such that $-90 \leq i \leq 90$ and $-180 \leq j \leq 180$.

3.2.2 *Determining start time t_0*

Secondly, miners receive a start time point t_0. Since the ground truth values of environmental variables are made public with a 5-day delay, we provide a dynamic approach to determine start time point t_0. Therefore, t_0 is constructed by calculating the difference between the time at the moment of sampling $t_{\text{GMT+0}}$ and a time offset p. The time $t_{\text{GMT+0}}$ follows the Copernicus time standard and the offset p is drawn from a bounded half-normal distribution:

$$t_0 = t_{\text{GMT+0}} + p$$

where

$$p = a + |X| \quad \text{with} \quad X \sim \mathcal{N}(\mu, \sigma^2)$$

Here, X follows a normal distribution with mean $\mu = 0$ and variance $\sigma^2 = 40$. The offset p is then obtained by taking the absolute value of X, shifting it by $a = -119$ hours and setting bounding constraints to ensure $-119 \leq p \leq 0$. This distribution ensures that most values of p are close to 120 hours (5 days) in the past, while allowing for more recent forecast initialization. This allows validators to evaluate miners faster, minimizing storage of forecast values and allow new miners to contribute. This design simulates the flexibility of dynamic time initialization, essential for real-world applications.

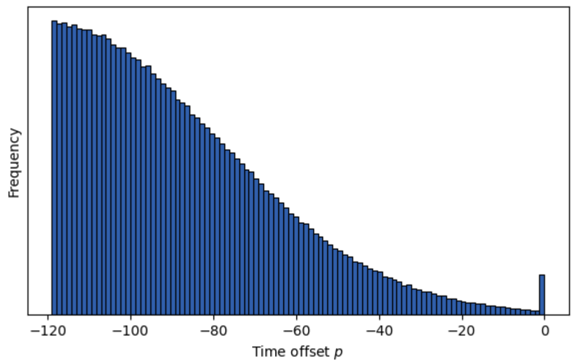

Fig. 1. Bounded half-normal distribution for sampling offset p

3.2.3 Sampling the forecast window h

Lastly, miners also receive a forecast window h. The forecast window h represents the number of steps ahead that miners must predict. It is sampled from a discrete uniform distribution

$$h \sim U\{a, b\}$$

where h is a random variable with lower bound $a = 1$ and upper bound $b = 24$, ensuring that forecasts extend between 1 and 24 hours from start time t_0.

3.2.4 Objective and Response Format

Given matrix A, start time t_0 and forecast window h, miners are challenged to perform an h-step ahead forecast $T+1, T+2, \ldots, T+h$ for every set of coordinates (i, j) denoted by:

$$\{\hat{Y}_{T+t,ij}\}_{t=1}^{h} = \{\hat{Y}_{T+1,ij}, \hat{Y}_{T+2,ij}, \ldots, \hat{Y}_{T+h,ij}\} \ \ \forall \ \ (i, j)$$

In addition, miners are tasked with constructing matrix B by replacing the respective elements in A with their corresponding time series of forecast values $\{\hat{Y}_{T+t,ij}\}_{t=1}^{h}$. For notational convenience, we simplify $\{\hat{Y}_{T+t,ij}\}_{t=1}^{h}$ to $\{\hat{Y}_{t,ij}\}$ by defining the indices of the forecast window h as $t = 1, \ldots, h$. Then, B is an $m \times n$ matrix denoted by:

$$B = \begin{bmatrix} \{\hat{Y}_{t,i_1 j_1}\} & \{\hat{Y}_{t,i_1 j_2}\} & \cdots & \{\hat{Y}_{t,i_1 j_n}\} \\ \{\hat{Y}_{t,i_2 j_1}\} & \{\hat{Y}_{t,i_2 j_2}\} & \cdots & \{\hat{Y}_{t,i_2 j_n}\} \\ \vdots & \vdots & \ddots & \vdots \\ \{\hat{Y}_{t,i_m j_1}\} & \{\hat{Y}_{t,i_m j_2}\} & \cdots & \{\hat{Y}_{t,i_m j_n}\} \end{bmatrix}$$

Matrix B shows a mapping between each unique set of spatial coordinates (i, j) and its corresponding unique set of matrix indices mn. Moreover, each value in element $\{\hat{Y}_{t,i_m j_n}\}$ represents the forecast value of temperature T_{2m} corresponding to the coordinate (i, j) at the forecast hour t.

3.3 Incentive mechanism

We designed a modular incentive mechanism ensuring fair assessment of each submission while allowing for specific needs to evolve the framework. At its core, the optimization goal is to minimize the average error between the miner prediction and the actual measured environmental values. An intuitive approach to measure this error is through the Root Mean Squared Error (RMSE), which penalizes larger deviations more heavily. However, RMSE alone does not incorporate the difficulty of the different forecasting challenges. To address this limitation, we introduce a weighted RMSE reward function, which adjusts the standard RMSE metric by incorporating a weighting factor that reflects the relative difficulty of each challenge.

3.3.1 Reward function

The reward function we use to evaluate miner performance is defined as:

$$\text{Reward} = \max\left(0, 1 - \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{\hat{y}_i - y_i}{z_i^*}\right)^2}\right)$$

Where:

- N: The number of forecast values in the response, capturing unique sets of time and location.
- $\hat{y}_i$: The forecasted value.
- y_i: The ground truth value.
- z_i^*: The transformed difficulty measure, denoted by:

  $$z_i^* = z_i * \text{scaling} + \text{offset}$$

  where the scaling and offset are configured by the validator.
- z_i: The difficulty measure.

The reward function is structured as follows:

- When a miner's forecast is highly accurate, the prediction error approaches zero, resulting in a near-zero weighted RMSE and a reward value close to 1.
- As the weighted RMSE increases, the reward decreases proportionally.
- If the unweighted RMSE exceeds z on average, the reward becomes 0, incentivizing miners to improve their predictions.

3.3.2 Ground-truth y_i

The ground truth values, denoted as y_i, represent the actual environmental measurements against which miner forecasts are evaluated. The data source for these ground-truth values is the ERA5 reanalysis dataset, provided by the Climate Data Store (CDS) under the European Union's Earth observation programme, Copernicus. ERA5 reanalysis integrates observational data from satellites, weather stations, and other sources using advanced numerical modeling techniques. This dataset offers hourly global measurements across a wide

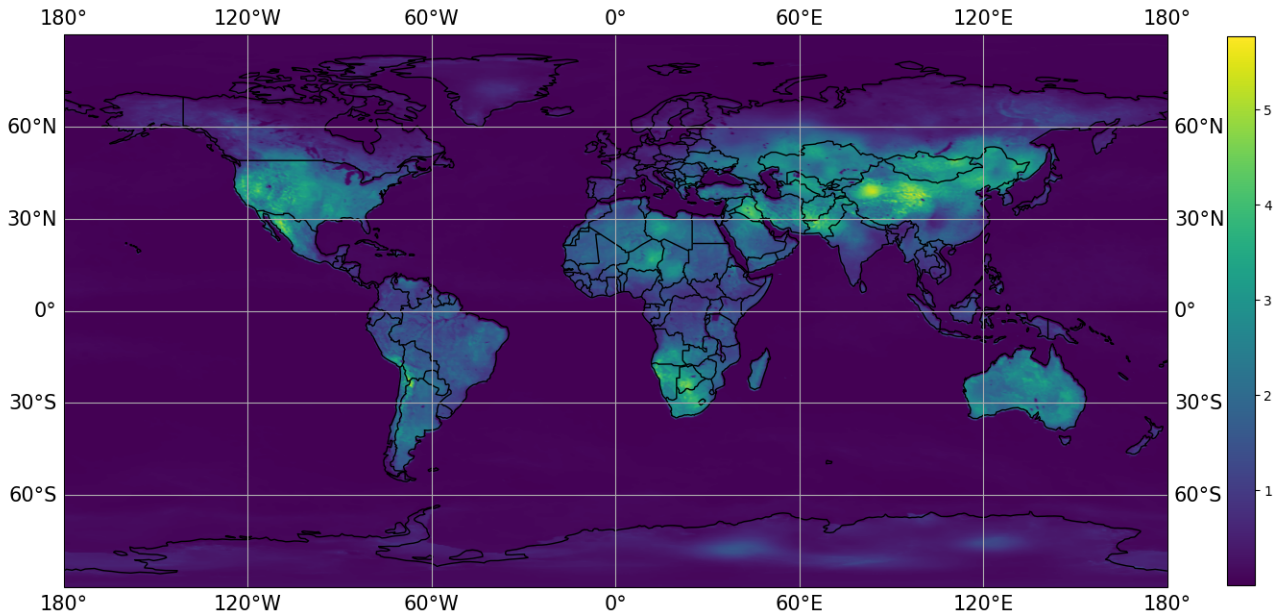

Fig. 2. Variance T_{2M} derivative Oct 2024

range of meteorological and environmental variables, making it the most comprehensive and reliable dataset for evaluating forecasting models.

3.3.3 Difficulty z_i

Climate forecasting difficulty varies significantly depending on geographical location. For instance, regions in the middle of the Pacific Ocean experience relatively stable temperature patterns, meaning a simple average would already obtain very low RMSE scores. Alternatively, mountainous regions exhibit high variability, making accurate predictions more challenging even for advanced models. Without an adaptive difficulty adjustment, miners could implicitly be penalized based on the random region assigned to them rather than their forecasting ability. To mitigate this bias, we introduce a difficulty-adjusted incentive mechanism using z_i.

The z_i value, or *difficulty* value plays a critical role in establishing an adaptive error threshold that accounts for temporal and spatial variability. A higher z_i value increases the tolerance for prediction error, whereas a lower z_i decreases the acceptable error margin.

The difficulty factor z_i is derived from the natural variability of environmental variables in different geographical locations. The underlying assumption is that regions with highly volatile temperature patterns are harder to forecast than regions with stable conditions. To quantify this, we compute the variance of temperature changes over time at each geographical location. This variance is calculated on a monthly basis to account for localized, short-term variability while avoiding the influence of seasonal trends or long-term climate shifts. The spatial distribution of temperature T_{2m} variance is visualized in Figure 2.

By basing the values of z on the natural variance in temperature fluctuations, this approach ensures that the forecast difficulty is data driven and objectively incorporated.

3.4 Hardware constraints for validators

Validation on Zeus is designed to be lightweight, aimed to minimise storage and hardware requirements. In addition, data processing is done locally, but since this has been highly optimised, validators will not need any GPU or CUDA support. Validators will only need a decent CPU machine, where we recommend having at least 8GB of RAM. Since data is loaded over the internet, it is useful to have at least a moderately decent (>3MB/s) internet connection.

4 Conclusion

This whitepaper introduces Zeus, a climate forecasting framework powered by Bittensor's blockchain technology. Through its decentralized architecture, Zeus fosters a competitive environment that drives continuous innovation, advancing the field of climate forecasting.

Zeus prioritizes accessibility, incentivizing engineers to develop and optimize forecasting models, and minimizing storage and hardware requirements for validators. This makes participation in the network easier and more inclusive. The decentralized architecture, ensures a global, distributed network of engineers who compete, fostering continuous innovation and preventing centralized control.

Out-of-sample forecasting makes the subnet ungameable by preventing miners from exploiting past data, while the z value further ensures fairness by adjusting the difficulty of forecast challenges based on historical temperature variability. Together, these mechanisms guarantee reliable and fair evaluations, regardless of geographic complexity.

These features make Zeus an accessible and decentralized framework for climate forecasting, set to revolutionize the field while ensuring fair rewards for every contribution along the way.

5 Future Work

In future iterations, Zeus will continue to scale and adapt to meet the dynamic needs of the climate forecasting landscape. Designed with flexibility in mind, the subnet can easily integrate new environmental variables in response to emerging technological advancements or real-world demands. This scalability enables Zeus to evolve into a more comprehensive forecasting framework, offering increasingly accurate and accessible forecasts while maintaining its competitive edge.

As the weather forecasting services market is projected to grow significantly, Zeus is positioned to provide substantial value by decentralizing the traditionally centralized and costly forecasting process. The subnet's adaptability opens opportunities for a wide range of applications, from disaster forecasting to grid management, ensuring its relevance across various industries. Furthermore, by incorporating flexible grid sizes and adjusting to validator needs, Zeus will enhance its ability to deliver timely, reliable, and market-driven forecasts. This approach not only increases the real-world impact of the system but also lays the foundation for monetization strategies, enabling Zeus to contribute meaningfully to the climate forecasting sector and capitalize on its broad utility.

References

[1] Vaughan, A., Markou, S., Tebbutt, W., Requeima, J., Bruinsma, W. P., Andersson, T. R., Herzog, M., Lane, N. D., Hosking, J. S., & Turner, R. E. (2024). Aardvark Weather: end-to-end data-driven weather forecasting. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2404.00411

[2] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., & Battaglia, P. (2023). Learning skillful medium-range global weather forecasting. *Science*, *382*(6677), 1416–1421. https://doi.org/10.1126/science.adi2336

[3] Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., Lam, R., & Willson, M. (2024). Probabilistic weather forecasting with machine learning. *Nature*. https://doi.org/10.1038/s41586-024-08252-9

[4] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2022). Pangu-Weather: a 3D High-Resolution model for fast and accurate global weather forecast. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2211.02556

[5] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., & Anandkumar, A. (2022). FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2202.11214