

# ORBIT-2: Scaling Exascale Vision Foundation Models for Weather and Climate Downscaling

Xiao Wang\*, Jong-Youl Choi\*, Takuya Kurihaya\*, Isaac Lyngaas\*, Hong-Jun Yoon\*, Xi Xiao<sup>§</sup>, Ming Fan\*  
Nasik Muhammad Nafi\*, Aristeidis Tsaris\*, Ashwin M. Aji<sup>†</sup>, Maliha Hossain\*, Mohamed Wahib<sup>‡</sup>  
Dali Wang\*, Peter Thornton\*, Prasanna Balaprakash\*, Moetasim Ashfaq\*, Dan Lu\*

\*Oak Ridge National Laboratory, Oak Ridge, United States

{wangx2, choij, kurihanat, lyngaasir, yoonh, fanm, nafinm, tsarisa, wangd, thorntonpe, pbalapra, mashfaq, lud1}@ornl.gov

<sup>†</sup>AMD Research and Advanced Development, Santa Clara, United States, ashwin.aji@amd.com

<sup>‡</sup>RIKEN Center for Computational Science, Kobe, Japan, mohamed.attia@riken.jp

<sup>§</sup>University of Alabama at Birmingham, xxiao@uab.edu

**Abstract**—Sparse observations and coarse-resolution climate models limit effective regional decision-making, underscoring the need for robust downscaling. However, existing AI methods struggle with generalization across variables and geographies and are constrained by the quadratic complexity of Vision Transformer (ViT) self-attention. We introduce ORBIT-2, a scalable foundation model for global, hyper-resolution climate downscaling. ORBIT-2 incorporates two key innovations: (1) Residual Slim ViT (Reslim), a lightweight architecture with residual learning and Bayesian regularization for efficient, robust prediction; and (2) TILES, a tile-wise sequence scaling algorithm that reduces self-attention complexity from quadratic to linear, enabling long-sequence processing and massive parallelism. ORBIT-2 scales to 10 billion parameters across 32,768 GPUs, achieving up to 1.8 ExaFLOPS sustained throughput and 92–98% strong scaling efficiency. It supports downscaling to 0.9 km global resolution and processes sequences up to 4.2 billion tokens. On 7 km resolution benchmarks, ORBIT-2 achieves high accuracy with  $R^2$  scores in range of 0.98–0.99 against observation data.

## I. PROBLEM OVERVIEW

Many regions lack dense ground-based observational networks, hindering early warning systems, disaster risk mitigation, and climate adaptation planning. In such cases, global climate models provide a crucial alternative, simulating atmospheric processes at planetary scale. However, their coarse resolution limits the representation of fine-scale phenomena, constraining accuracy at regional levels.

Downscaling bridges this gap by translating coarse resolution global climate model outputs into fine-scale outputs [1]. This process is critical across a wide range of sectors, including agriculture [2], water resources [3], infrastructure and urban planning [4], energy systems [5], and extreme event forecasting [6]. Despite its importance, downscaling remains both scientifically and computationally challenging. It requires physically consistent predictions from massive, high-dimensional spatiotemporal data while maintaining accuracy across diverse regions. These challenges highlight the urgent need for scalable, high-fidelity downscaling approaches to enable effective climate services, disaster preparedness, and policy planning.

Traditional downscaling methods fall into two main categories: dynamical [7], which uses nested physical models to simulate fine-scale processes but is computationally intensive and limited to regional domains; and statistical [8], [9], which is computationally efficient but often lacks physical fidelity

and generalizability. More recently, artificial intelligence (AI) has emerged as a powerful alternative, offering the potential for high-resolution predictions with significantly lower inference costs [10], [11], [12]. Task-specific deep learning models have been developed to learn mappings between coarse and fine-resolution fields, typically falling into two categories: non-generative models [13], which are efficient and stable but may lack fine spatial detail, and generative models, such as diffusion-based approaches [14], [15], [16], which produce sharper outputs at the expense of much higher computational and training costs. Despite promising results, most task-specific models must be retrained for each variable, resolution, or region and often struggle to generalize across diverse, physically distinct climate variables [17].

To overcome the limitations of task-specific models, recent efforts have introduced foundation models such as Prithvi [18] and ClimateLearn [19], which employ multi-task Vision Transformer (ViT)-based architectures to support downscaling across variables and geographic regions. While these models represent an important step forward, they remain constrained by resolution limits, computational cost, and model scalability.

A primary bottleneck is the computational complexity of downscaling at high resolution. For instance, Prithvi achieves 12 km resolution over Europe but is restricted to 50–60 km globally due to the quadratic scaling of ViT self-attention [20]. ViTs divide spatial data into patches, treating each patch as a token. As resolution increases, the number of tokens grows, and self-attention computes pairwise interactions among all tokens, resulting in quadratic growth in memory and compute demands. Unlike Natural Language Processing (NLP) models, which operate on one-dimensional text sequences and scale to over one million tokens [21], ViTs handle high-dimensional spatial inputs with complex dependencies across multiple axes, making long-sequence scaling significantly more computationally intensive. As a result, the longest ViT sequence reported to date is limited to 188k tokens [22]. This constraint directly limits the maximum data resolution ViTs can process, as sequence length scales proportionally with spatial resolution.

Another significant challenge is the uncertainty associated with multi-variable downscaling, as translating coarse-resolution data to fine scales is a highly ill-posed inverse problem—one that becomes even more complex when multiple

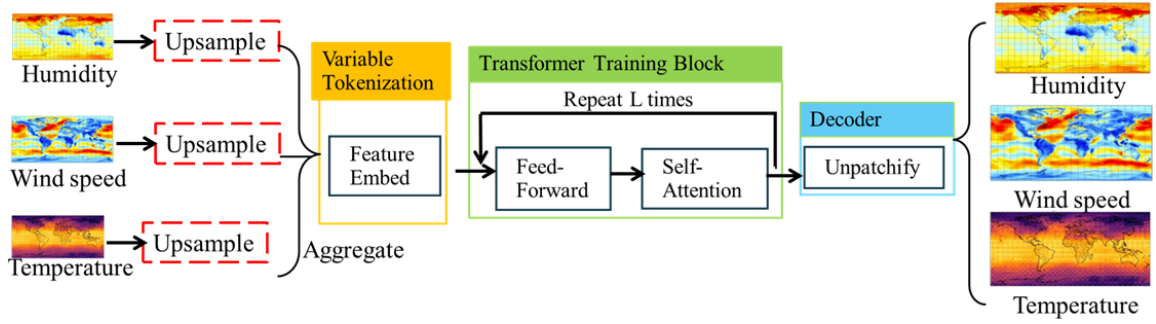


Fig. 1: A generalized AI architecture diagram for state-of-the-art downscaling foundation models. Note that upsampling is used for each channel prior to training blocks to reduce downscaling uncertainty.

climate variables are involved. Unlike super-resolution tasks in computer vision [23], [24], where Red-Green-Blue channels represent the same physical quantity, climate variables such as temperature, humidity, and wind are governed by distinct yet interrelated physical processes. This heterogeneity increases the difficulty of learning consistent mappings and exacerbates uncertainty in predictions. A common mitigation strategy is to upsample coarse inputs prior to training [18], [19], which can help reduce uncertainty but significantly increases sequence length and, in turn, computational cost due to ViT’s quadratic complexity. Moreover, upsampling introduces artifacts that can propagate through the model, limiting its effectiveness.

One further limitation of existing downscaling foundation models is their restricted model scale. For example, Prithvi [18] is constrained to 1.4 billion parameters, primarily due to the computational difficulty of scaling ViTs for high-dimensional spatiotemporal data. A major advancement in this area is the Oak Ridge Base AI foundation model for Earth System Predictability (ORBIT) [17], which leverages hybrid sharding and orthogonal parallelisms to scale ViTs to 113 billion parameters—five times larger than previous ViTs and more than 100× larger than typical climate models. While ORBIT represents a major milestone in large-scale Earth system modeling, it is specifically designed for temporal forecasting and does not address spatial downscaling. In particular, it does not resolve the ViT long-sequence bottleneck nor mitigate the uncertainty associated with inverse downscaling problems.

To address the limitations of current foundation models in high-resolution downscaling and to extend the ORBIT framework, we introduce ORBIT-2, a scalable and computationally efficient foundation model for climate downscaling. At its core is a novel ViT architecture, *Residual Slim ViT (Reslim)*, specifically designed to bypass the high computational cost associated with traditional upsampling-based approaches. Unlike existing models that upsample inputs to mitigate uncertainty—resulting in quadratic increases in memory and computation—Reslim operates directly on adaptively compressed spatial inputs, significantly reducing sequence length while preserving critical information. It preserves accuracy and reduces uncertainty through a lightweight residual learning architecture, enabling efficient, low-overhead predictions. Additionally, both training and inference are framed as a *Bayesian Estimation* problem, incorporating a Markov Random Field

Total Variation prior to further constrain uncertainty and improve spatial consistency.

Complementing this architecture is the *Tile-Wise Sequence Scaling Algorithm (TILES)* that reduces ViT’s self-attention complexity from quadratic to linear. It works by dividing images into overlapping tiles, each processed in parallel on separate Graphical Process Units (GPUs) using localized self-attention. Each tile’s downscaled outputs are then seamlessly merged to the full image. This strategy enables efficient and scalable ViT-based downscaling, making ultra-high-resolution, global-scale applications computationally feasible.

Leveraging the above innovations, ORBIT-2 sets a new benchmark for AI-driven climate and weather downscaling through four key breakthroughs:

- **Efficient Reslim Architecture** by operating directly on compressed inputs, achieving over 660× speedup compared to standard ViTs—without compromising accuracy.
- **Longest ViT Sequence Length** by scaling ViT sequence lengths to unprecedented levels—up to 4.2 billion tokens for a 9.5M parameter model and 671 million tokens for a 10B model—surpassing the prior state-of-the-art of 188K tokens by several orders of magnitude [22]. This eliminates the long-standing sequence bottleneck, enabling global downscaling at resolutions as fine as 0.9 Kilometer (km).
- **Scalable Large Model Training** by training models with up to 10 billion parameters across 32,768 GPUs, achieving 92–98% strong scaling efficiency and sustained throughput of up to 1.8 ExaFLOPS.
- **State-of-the-Art Accuracy** achieving  $R^2$  scores of 0.98 for precipitation and 0.99 for temperature at 7 km resolution over the continental United States, setting a new standard in high-fidelity downscaling performance.

## II. BACKGROUND & STATE OF THE ART

Figure 1 illustrates the generalized architecture of leading downscaling foundation models, including Prithvi [18] and ClimateLearn [19]. The inputs consist of low-resolution data with multiple atmospheric physical variables, normalized and bias corrected, and each channel of the architecture reads data for a distinct variable. To address downscaling inverse problem uncertainty, current models upsample coarse-resolution inputs, either via interpolation [19] or convolution [18], before training. This upsampling process is crucial, as it provides

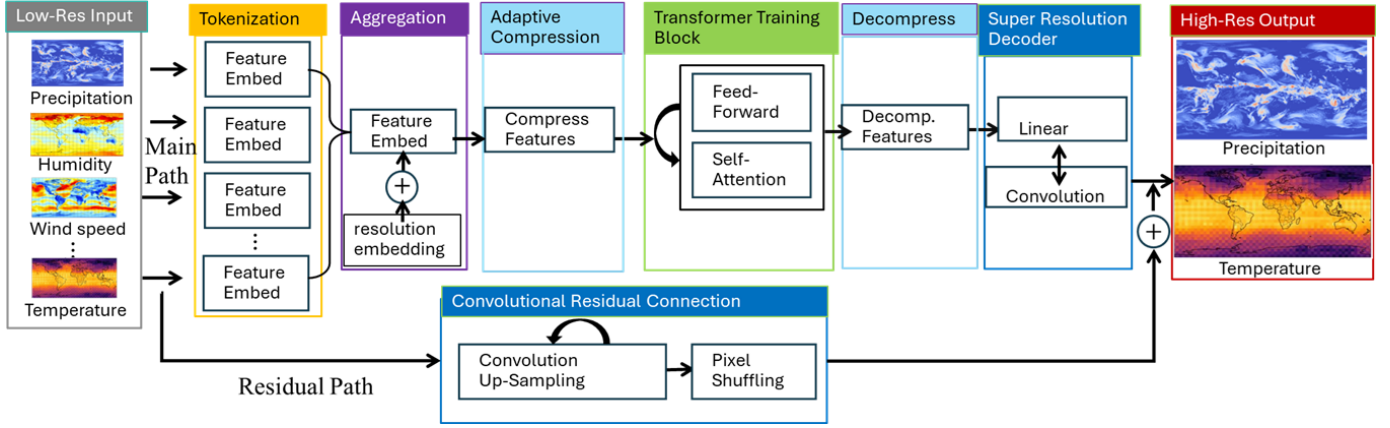


Fig. 2: Reslim architecture is split into main and residual paths. No upsampling is used for the main path for ViT training, leading to reduced computations. Residual path is used to condition prediction for reduced uncertainty.

a higher-resolution baseline for ViT training, mitigating uncertainty from the inherently ill-posed nature of the multi-variable downscaling problem, thereby improving accuracy and uncertainty. Once upsampled, multi-channel inputs are aggregated into a single-channel representation in feature space, a step that can be performed using either cross-attention mechanisms [25] or shallow convolutional layers [18], [19]. This aggregated representation is then trained by the ViT training blocks, consisting of self-attention and feedforward sub-layers. Finally, the trained output is projected back from feature to image space for each individual physical variable.

This approach, however, introduces major challenges. Upsampling coarse-resolution input data before training increases the sequence length, which increases in proportion to the resolution increase, causing a quadratic increase in memory and computations due to ViT’s self-attention mechanism. This severely limits scalability and resolution, leaving the long-sequence bottleneck unresolved. Prithvi, for example, is limited to relatively coarse 50-60 km resolution for global downscaling. To address this, prior work proposed both AI architecture and High Performance Computing solutions.

**Architecture solutions.** To mitigate this, architectures like Swin Transformer alleviate some of the computational burden by introducing a hierarchical architecture with shifted window attention [26], [27]. Instead of processing the entire image at once, Swin Transformer partitions the image into smaller, non-overlapping local windows, where self-attention is computed independently within each window. To capture global spatial dependencies, features learned from local windows are aggregated into global features through an architecture hierarchy. While this reduces computing complexity, Swin Transformer has fundamental limitations and its layers of architecture hierarchy must scale proportionally with higher resolution, making it unsuitable for foundation models that needs a single model to generalize across diverse datasets with varying resolutions. Additionally, Swin Transformer’s model size grows with the architecture hierarchy, shifting the computational bottleneck from long-sequence processing to large-model scaling. Consequently, Swin Transformer can only scale up to 147K sequence length on standard 3-channel

images [27], far below what is needed for high-resolution, multi-variable downscaling.

Other sparse attention architectures, such as MaxViT [28], attempt to mitigate computational cost by sampling self-attention computations. While this reduces complexity, it comes at the expense of accuracy degradation when the sampling ratio is too high, and it does not address the fundamental quadratic complexity long-sequence problem.

**Scaling algorithm solutions.** Besides architecture innovations, scaling algorithms, such as sequence parallelism [22], [29], [30], has been proposed as an alternative strategy for scaling ViT sequence length. It distributes image patch tokens across GPUs for parallel computing, alleviating memory constraints. However, because self-attention requires each token to interact with all other tokens from every other GPU, sequence parallelism incurs substantial inter-GPU communication overhead and limits its scalability. More critically, it does not resolve the fundamental quadratic complexity, which causes computational costs to grow rapidly with increased downscaling resolution. As a result, current ViT sequence parallelisms are limited to a maximum of 188K token sequence lengths [22], which remain insufficient for high-resolution multi-variable downscaling.

It is also important to note that other commonly used parallelisms—such as Fully Sharded Data Parallelism (FSDP) [31], Tensor [32], pipeline [33], [34], [35] and hybrid sharded parallelisms [17] are all designed to scale model sizes, rather than long sequences of high-resolution and high-dimensional spatial data. Consequently, none of the existing model parallelisms fundamentally overcome the long-sequence bottleneck in ViTs required for high-resolution global downscaling and there is an urgent need to develop computing efficient and massively parallel architecture and scaling algorithm.

### III. INNOVATION REALIZED

#### A. Reslim: A Lightweight ViT Architecture for Scalable and Uncertainty-Aware Downscaling

Unlike existing foundation models that rely on input up-sampling to establish downscaling baselines, which leads to increased sequence length and high computational cost,

ORBIT-2 introduces Residual Slim ViT (Reslim), a highly efficient architecture that significantly reduces training time and memory usage without compromising accuracy. The key innovation of Reslim is its ability to operate directly on low-resolution and adaptively compressed inputs, drastically reducing sequence length and computational burden. To counteract the uncertainty typically introduced by bypassing upsampling prior to ViT training, Reslim incorporates Bayesian estimation and a residual convolutional learning path, enabling high accuracy while maintaining efficiency. Its non-hierarchical design further promotes generalization across datasets with varying spatial resolutions, making it well-suited for scalable, foundation-level Earth system modeling.

**Main ViT Path.** Figure 2 illustrates the Reslim architecture. After tokenizing each low-resolution physical variable into feature embeddings, the model proceeds along two architectural paths: the main ViT and residual paths. Crucially, the main path eliminates input upsampling, avoiding the sequence length inflation and the quadratically increased computing cost typical of ViT architectures.

First, the main path uses a cross-attention module to aggregate multi-variable embeddings into a unified representation, effectively collapsing the variable dimension. A learnable resolution embedding encodes the desired output resolution and is added to the feature embedding, enabling resolution-aware predictions—an essential capability for modeling resolution-dependent Earth system behaviors. Next, an optional adaptive spatial compression module, which will be explained further in the next paragraph, reduces the sizes of the embeddings before they are passed through ViT training blocks. When enabled, this module compresses spatial features; otherwise, it acts as an identity function. After processing, a decoder comprising convolutional layers and linear projections reconstructs the high-resolution output.

**Adaptive Spatial Compression.** Our objective is not only to train directly on low-resolution inputs, but also to further reduce token count and computational cost through compression. Reslim achieves this via an adaptive spatial compression technique, inspired by adaptive image patching and mesh refinement methods [36]. After aggregating multi-variable features (purple block in Fig. 2), the model projects the embedding back into image space and recursively partitions it into spatial quadrants using a quad-tree structure. Partitioning continues for any quadrant where the estimated feature density—computed via Canny edge detection—exceeds a predefined threshold, terminating when a minimum patch size is reached or below predefined threshold.

This approach enables finer-grained learning in feature-rich regions through smaller patches, and coarse-grained learning to smoother regions through larger patches, where less detail is needed. Figure 3 illustrates an example image after variable-aggregated features are mapped back to image space. Compared to conventional uniform patching (Fig.3(a)), where each grid represents an image patch token, the adaptive spatial compression method (Fig.3(b)) reduces the number of patch tokens by 7x in this figure example, significantly decreasing sequence length and computing cost. After ViT training blocks, the decompression module reconstructs the high-resolution output from the compressed embeddings.




Fig. 3: Comparison with and without adaptive spatial compression. Each yellow grid is an image patch.

**Residual Learning.** Reslim improves computational efficiency by removing the upsampling step from the main ViT path and training directly on low-resolution, spatially compressed inputs. This design dramatically shortens sequence lengths and reduces the quadratic computational cost typically associated with ViT training. However, bypassing input upsampling introduces uncertainty, as conventional foundation models rely on upsampled inputs to provide a coarse downscaling baseline. Reslim addresses this challenge through two complementary innovations: residual convolutional learning and a Bayesian estimation objective.

The residual convolutional path reintroduces upsampling outside the main ViT path, using lightweight convolutional layers with linear complexity. This path generates a high-resolution approximation that is added to the ViT output before loss computation. Such design yields two major benefits: (1) it avoids the expensive quadratic cost of increasing the ViT sequence length due to upsampling. The upsampling is moved to the residual path, where convolutional layers have linear complexity to input size and thereby upsampling in the residual path incurs minimal computing cost. (2) it simplifies the learning task by letting the ViT focus on predicting the residual difference between the convolutional approximation and the ground truth, rather than the full downscaling transformation. This soft constraint stabilizes training, enhances physical plausibility, and significantly reduces downscaling uncertainty. As a result, Reslim achieves high downscaling accuracy with significantly reduced computations compared to conventional ViT.

**Bayesian Training Loss.** To further reduce uncertainty and improve accuracy, Reslim reformulates its training as the following Bayesian optimization problem with a Generalized Markov Random Field Total Variation prior:

$$\hat{x} \leftarrow \underset{\hat{x}}{\operatorname{argmin}} \|y - \hat{x}\|_D^2 + \sum_{k=1}^K \sum_{i=1}^N \sum_{j \in C(\hat{x}_{k,i})} b_{i,j} \|x_{k,i} - x_{k,j}\| ,$$

where  $y$  is the high-resolution ground-truth,  $\hat{x}$  is the Reslim prediction, and  $D$  is a latitude weighting matrix to account for the decrease in longitudinal spacing toward the poles.  $K$  is the number of output variables, and  $N$  is the total number of pixels per variable. The neighborhood  $C(\hat{x}_{k,i})$  contains all spatial neighbors of pixel  $\hat{x}_{k,i}$ , which is the  $i^{\text{th}}$  pixel for the  $k^{\text{th}}$  variable.  $b_{i,j}$  is a spatial weighting factor inversely proportional to the euclidean distance between each pixel pair in the same neighborhood. In the above formulation, the first term,  $\|y - \hat{x}\|_D^2$ , is the Bayesian forward data likelihood term using a latitude-weighted mean squared error. The second term



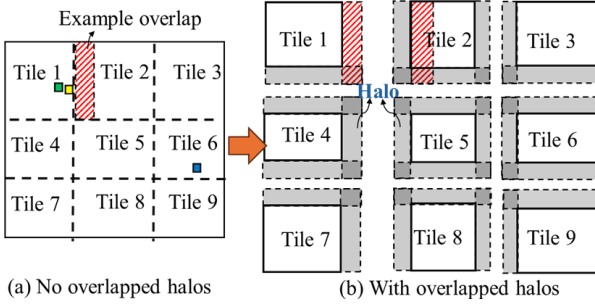


Fig. 4: (a) TILES divides inputs and outputs into tiles, and each GPU downscales a tile separately. (b) Overlapped halos are added to each tile to address border artifacts among neighboring tiles.

is a total variation spatial prior, promoting local smoothness by penalizing irregularities within local neighborhoods, but also preserving edges and discontinuities. This makes it well suited for downscaling tasks that depend on spatial coherence and structure preservation.

#### B. TILES: Tilewise Efficient Sequence Scaling Algorithm

While Reslim significantly reduces computation by operating on low-resolution and compressed inputs, it does not resolve the inherent quadratic complexity of self-attention. As resolution increases, this limitation becomes a bottleneck. To address this, we introduce the Tilewise Efficient Sequence Scaling Algorithm (TILES), a scalable sequence processing strategy that reduces attention complexity from quadratic to linear and enables efficient parallelization across GPUs.

TILES is motivated by the spatial locality property of downscaling, where the downscaling for each high-resolution pixel is primarily influenced by spatially nearby coarse-resolution inputs. This “point spread” effect, well studied in the remote sensing literature [37], [38], [39], implies that long-range pixel correlations can be safely ignored without affecting downscaling accuracy. For example, neighboring pixels such as the green and yellow ones in Fig. 4(a) have high mutual influence, while distant pairs (e.g., green and blue) contribute minimally or no influence to each other’s predictions.

Leveraging this locality, TILES partitions each input and downscaling output into spatial tiles, assigning each tile to a separate GPU. Each GPU then performs downscaling separately for its assigned tile, and self-attention is restricted within each tile, preserving local context while ignoring long-range dependencies across tiles. This tilewise downscaling reduces self-attention complexity from quadratic to linear. More specifically, the computation complexity is  $O(\frac{N^2}{T})$ , where  $N$  is the number of image patch tokens and  $T$  is the number of tiles. For fixed-size tiles,  $T$  increases proportionally with  $N$ , making the overall complexity linear.

However, strict tiling introduces border artifacts, as pixels near the border of each tile lack context from neighboring tiles. To mitigate this, TILES introduces halo padding, where each tile is padded with a fixed-width halo (shown in gray in Fig.4(b)) that overlaps adjacent tiles. For instance, the red region in Fig.4(a) is shared between Tile 1’s halo and Tile




Fig. 5: Orthogonal levels of parallelisms mapped to supercomputer at cluster, node, and device levels.

2’s core region both shaded in red in Fig.4(b). This overlap restores spatial continuity across tile boundaries, ensuring that border pixels—such as the yellow pixel—receive complete local context. Note that the halo width is determined empirically. Larger halos improve accuracy but increase computation, while smaller halos reduce cost but risk accuracy loss.

After each GPU independently downsamples its tile, the halo regions are discarded, and the non-padded tile outputs are stitched together to form the final high-resolution output. Since each GPU processes a different tile, leading to different gradient and model parameter update, gradients from all GPUs are averaged to maintain the model consistency across GPUs and global optimization. This inter-GPU communications, however, have minimal communication frequency and overhead as it takes place only once per data batch.

#### C. Orthogonal Parallelisms

TILES efficiently scales sequence length, enabling high-resolution downscaling. However, it does not address model size scaling. To support both large foundation models and global high-resolution downscaling, TILES must be integrated with complementary model-parallel strategies. Since TILES and model parallelism target orthogonal goals, with TILES for sequence length and model parallelisms for model size, they can be combined seamlessly. This results in a unified framework incorporating four distinct parallelism strategies:

- TILES sequence parallelism: Distributes long sequence lengths for ViTs for tilewise approximation as discussed before. Requires least communication overhead.
- Fully Sharded Data Parallelism (FSDP) [31]: shards both data and model parameters across GPUs, but requires temporarily gathering the full model during forward and backward passes. Requires moderate communication overhead.
- Tensor Model Parallelism [32]: Only shards model parameters and keep parameters sharded throughout training. Requires most communication overhead.
- Distributed Data Parallelism (DDP) [40]: distributes only the training data without sharding model parameters. Requires least communication overhead.

Fig. 5 shows how these orthogonal parallelisms map to a supercomputer hardware. Two adjacent nodes form a TILES sequence parallel group (green dashed boxes), responsible for

Pretraining							
Dataset name	Region	Resolution (km)	Input Vars	Output Vars	Sample Size (in/out)	# Sample Pairs	Size (GB)
ERA5 $\rightarrow$ ERA5	Global	622 $\rightarrow$ 156	23	3	[32, 64, 23] $\rightarrow$ [128, 256, 3]	367,920	200
ERA5 $\rightarrow$ ERA5	Global	112 $\rightarrow$ 28	23	3	[180, 360, 23] $\rightarrow$ [720, 1440, 3]	367,920	6,328
PRISM $\rightarrow$ PRISM	US	16 $\rightarrow$ 4	7	3	[180, 360, 7] $\rightarrow$ [720, 1440, 3]	14,235	189
DAYMET $\rightarrow$ DAYMET	US	16 $\rightarrow$ 4	7	3	[180, 360, 7] $\rightarrow$ [720, 1440, 3]	14,946	200
Fine-Tuning							
[ERA5, DAYMET] $\rightarrow$ DAYMET	US	28 $\rightarrow$ 7	23	3	[120, 240, 23] $\rightarrow$ [480, 960, 3]	14,946	113
Model Inference Evaluation							
ERA5 $\rightarrow$ IMERG	Global	28 $\rightarrow$ 7	23	3	[720, 1440, 23] $\rightarrow$ [2880, 5760, 3]	1,488	132

TABLE I: Datasets used for pretraining, fine-tuning, and inference. Each entry specifies the downscaling resolution, input/output variable counts, sample dimensions, number of training samples, and total storage size.

scaling sequence lengths. Multiple sequence parallel groups form a DDP group, distributing data batches across the system.

Within each sequence parallel group, GPUs participate in both tensor model and FSDP parallelisms for model scaling. Tensor model parallelism operates within a node, leveraging its low-latency interconnect to mitigate communication overhead for hidden dimension partitioning. FSDP (red dashed boxes) spans GPUs across neighboring nodes within the same TILES group, enabling parameter and data sharding.

Once all parallelisms are established, each GPU receives a subset of the model and data that can be optimized with Flash Attention [41] to reduce memory use and cache misses, as detailed in the next subsection. Within each GPU, streaming multiprocessors (SMs) are mapped to Flash Attention cache blocks, executing vector operations in parallel within each block. Meanwhile, CPUs asynchronously load data and construct quad-trees to track the spatial layout of adaptively compressed patches described in Fig. 3.

Note that this multi-level strategy aligns the parallelism hierarchy with the hardware architecture to optimize performance. Neither DDP nor TILES sequence parallelisms requires frequent communication, and are therefore mapped to cluster nodes with slower network communication. Tensor and FSDP model parallelisms requires more frequent communication, and are therefore mapped to GPUs within the same node and across neighboring nodes to utilize their faster in-node and neighboring-node network communications. The Flash Attention require the most frequent communication and are therefore mapped to SMs within the same GPU, which has the fastest network through shared L2 cache.

#### D. Optimizations

To further boost performance, we applied these optimizations:

**Hybrid-OP Parallelism.** We adopt the Hybrid-OP optimization technique from ORBIT [17], which leverages the mathematical structure of matrix chain multiplication to shard model parameters in alternating row and column dimensions. This optimization combines tensor model parallelism with FSDP, achieving superior scalability with reduced communication overhead and frequency compared to without Hybrid-OP.

**Flash Attention.** To accelerate self-attention computation, we use Flash Attention [41], which applies a cache-blocking technique to minimize memory access to GPU global memory. By maximizing data reuse from high-bandwidth on-chip cache,

Flash Attention significantly improves compute throughput through higher cache hit rates and faster memory access.

**Mixed Precision and Layer Wrapping.** We further utilize BFLOAT16 mixed-precision to speed up training while reducing memory usage. To address numerical instability—where gradients with extreme magnitudes may underflow or overflow in BFLOAT16—we apply PyTorch’s dynamic gradient scaling [42]. This technique automatically rescales gradients into a representable range and reverses the scaling during parameter updates, ensuring numerical stability.

To further reduce communication cost, we apply FSDP in a layer-wise fashion [17]. Instead of sharding all model layers in a single instance, parameters are sharded one layer at a time. This reduces synchronization overhead and memory use.

#### IV. HOW PERFORMANCE WAS MEASURED

**Model Configuration.** All experiments in Sec. V use four model configurations: 9.5M (256-dim embedding, 6 layers, 4 heads), 126M (1024-dim, 8 layers, 16 heads), 1B (3072-dim, 8 layers, 24 heads), and 10B (8192-dim, 11 layers, 32 heads) parameters.

**System Details.** Experiments were conducted on the Frontier supercomputer at Oak Ridge National Laboratory. Each node consists of one 64-core AMD EPYC CPU and 8 GPUs (64 GB each), organized into 4 MI250X cards with two GPUs per card. GPUs on the same card communicate via Infinity Fabric CPU-GPU, while all four MI250X cards are connected via 50 GB/s GPU-GPU Infinity Fabric. Nodes are interconnected using 100 GB/s Slingshot-11. The software stack includes PyTorch v2.6, ROCm v6.2.4, and libfabric v1.22.

**Datasets.** Table I summarizes the datasets used for pretraining, fine-tuning, and inference. Our model is trained on paired input  $\rightarrow$  output datasets designed to achieve 4X spatial refinement. For pretraining at global scale, we use the ERA5 reanalysis dataset [43] (1980–2020), employing two sets of resolution pairs: 622 km  $\rightarrow$  156 km and 112 km  $\rightarrow$  28 km. The data are split into 38 years for training, 2 years for validation, and 1 year for testing. The ERA5 dataset includes 23 variables: 5 static fields, 12 atmospheric variables (humidity, wind speed, and temperature at 200, 500, and 850 hPa), and 6 surface variables. For United States (US)-focused pretraining, we utilize the PRISM and DAYMET observation datasets (1980–2022) [44], [45], performing 4X downscaling from 16 km to 4 km.

Fine-tuning is performed over the US by using both ERA5 and DAYMET at 28 km as inputs and 7 km DAYMET

(a) Reslim Architecture Speedup Comparison with ViT									
Arch	Model Size	Resolution (km)	Seq. Length	Compression	Tiles	Time/sample (s)	Speedup	PSNR	SSIM
ViT	9.5M	622→156	24,576	1×	1	7.3e-4	1	35.0	0.94
Reslim	9.5M	622→156	24,576	1×	1	1.1e-6	660	36.7	0.96
ViT	9.5M	112→28	777,660	1×	1	OOM	NA	NA	NA
Reslim	9.5M	112→28	777,660	1×	1	1.2e-3	NA	37.6	0.96
(b) Adaptive Compression & Tiling Speedup Comparison with Reslim Baseline									
Reslim	9.5M	112→28	777,660	8×	1	3.6e-4	3.3	37.7	0.96
Reslim	9.5M	112→28	777,660	16×	1	1.8e-4	6.6	37.8	0.96
Reslim	9.5M	112→28	777,660	32×	1	1.7e-4	7.1	37.9	0.96
Reslim	9.5M	112→28	777,660	1×	4	8.0e-4	1.5	37.7	0.96
Reslim	9.5M	112→28	777,660	1×	16	6.3e-4	1.9	37.7	0.96
Reslim	9.5M	112→28	777,660	1×	36	7.4e-4	1.6	37.7	0.96

TABLE II: (a) Computing performance comparison between ViT and Reslim at 128 GPUs. For 622→156 km downscaling, Reslim achieves a 660× speedup over conventional ViT while maintaining accuracy. (b) Illustrates Reslim’s performance gains at varying adaptive compression rates and tile counts, relative to a Reslim baseline without compression or tiling.

Architecture	Model Size	Compression	Tiles	GPUs	Max Seq. Length	Output Size	Global Resolution (km)
ViT	9.5M	1×	1	8	25K	[128, 256, 18]	156
ViT	10B	1×	1	8	OOM	—	—
Reslim	9.5M	1×	1	8	298M	[5760, 11520, 18]	3.5
Reslim	9.5M	1×	1	32	466M	[7200, 14400, 18]	2.7
Reslim	9.5M	4×	16	8	1.1B	[11520, 23040, 18]	1.7
Reslim	9.5M	4×	16	128	4.2B	[21600, 43200, 18]	0.9
Reslim	10B	1×	1	8	18M	[1440, 2880, 18]	14
Reslim	10B	4×	16	8	74M	[2880, 5760, 18]	6.9
Reslim	10B	4×	16	512	671M	[8640, 17280, 18]	2.3

TABLE III: Maximum sequence length scaling across architectures, model sizes, compression, tiles and GPU count.

as output ground truth. The fine-tuning dataset is split into training, validation, and testing in the same way as pretraining.

After fine-tuning, we apply the trained model for inference to downscale global ERA5 precipitation data from 28 km to 7 km resolution, and evaluate against IMERG observation dataset [46], a NASA product for global precipitation observation. We compare the model’s daily aggregated 7 km predictions with IMERG globally. Since inference does not include bias correction—and both ERA5 and IMERG contain uncertainties—perfect alignment is not expected.

**Performance Metrics.** The total number of floating point operations (FLOPs) of the systems was collected via the Microsoft DeepSpeed Profiler [47] and we only gathered the FLOPs on GPUs. Only the mixed-precision BFLOAT16 results were reported. We measured following performance:

- *Time-to-solutions.* Defined as the average wall-clock time to downscale each hourly sample. Equivalent to epoch runtime divided by the total number of sample count. We reported numbers for both 622→156 and 112→28 km resolutions.
- *Strong scaling efficiency.* Measured speedup per epoch relative to GPU count, with the runtime at 512 GPUs (64 nodes) as the 100% baseline.
- *Sustained throughput.* It is defined as the total FLOPs to downscale each data point, divided by the average wall clock time in seconds. The performance includes the whole application with IO. Reported in ExaFLOPS.
- *Accuracy.* We use both scientific and image-based metrics for downscaling accuracy against observations: Coefficient of determination ( $R^2$ ), Root-Mean-Square-Error (RMSE), RMSE for Quantiles, Structural Similarity Index (SSIM), and Peak-Signal-Noise-Ratio (PSNR). Higher  $R^2$ , SSIM, and PSNR scores represent higher fidelity downscaling, while lower RMSE represents higher fidelity downscaling.

**Dataset, Source Code and Model Checkpoint Release.** We will publicly release the datasets, source code and trained model checkpoints soon and update this manuscript.

## V. PERFORMANCE RESULTS

### A. Speedup Ablation Studies

Table II(a) presents an ablation study comparing the performance of ViT and Reslim architectures, using models with 9.5M parameters and at 128 GPUs. Two ERA5→ERA5 downscaling tasks were used: 622→156 km and 112→28 km (details in Table I). Each hourly output sample is tokenized into image patches. For the 622→156 km task, outputs of shape [128, 256, 3] and 2×2 patch size yield sequence length of 24,576; for 112→28 km, output size [720, 1440, 3] produces 777,660 tokens. No adaptive spatial compression or tiling was applied in this comparison.

The seventh column of Table II(a) reports the average time to downscale each hourly sample. The eighth column shows the speedup from Reslim relative to the ViT baseline. Notably, the Reslim architecture avoids expensive upsampling operations by operating directly on low-resolution inputs, resulting in significant computational savings. For the smaller 622→156 km task, Reslim achieves a 660× speedup over ViT at the same number of GPUs while maintaining comparable accuracy, as measured by PSNR and SSIM. This demonstrates the effectiveness of Reslim’s residual learning design and Bayesian training loss in maintaining predictive accuracy while reducing computational cost. For the larger 112→28 km resolution task, the ViT model fails due to out-of-memory (OOM) errors. Consequently, a direct speedup comparison is not available, although Reslim completes the task efficiently and maintains high accuracy.

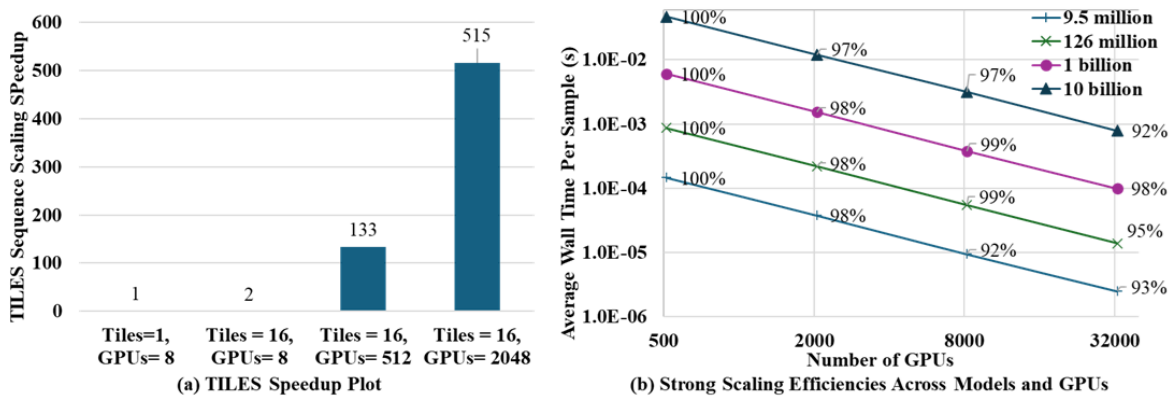


Fig. 6: (a) TILES sequence scaling algorithm speedup across GPUs, compared to an 8-GPU baseline that does not utilize tiling. (b) Strong scaling efficiencies up to 4096 nodes (32,768 GPUs) for various model sizes, maintaining a strong scaling efficiencies of 92-98% at 4096 nodes.

(a) Evaluation metrics for minimum temperature (Kelvin)							
Model Size	R <sup>2</sup>	RMSE	RMSE $\sigma_1 > 68\%$	RMSE $\sigma_2 > 95\%$	RMSE $\sigma_3 > 99.7\%$	SSIM	PSNR
9.5M	0.991	3.812	4.652	9.704	15.497	0.958	29.02
126M	<b>0.999</b>	<b>0.505</b>	<b>0.630</b>	<b>1.025</b>	<b>1.491</b>	<b>0.987</b>	<b>45.96</b>
(b) Evaluation metrics for total precipitation (millimeter/day)							
9.5M	0.975	0.146	0.166	0.344	0.449	0.931	29.03
126M	<b>0.979</b>	<b>0.135</b>	<b>0.154</b>	<b>0.296</b>	<b>0.365</b>	<b>0.932</b>	<b>30.20</b>

TABLE IV: Comparison of downscaling accuracy for temperature and precipitation over the U.S. using models with 9.5M and 126M parameters. Results highlight performance gains from increased model capacity.

Table II(b) explores further speedup gains from adaptive spatial compression and sequence tiling, compared to the Reslim baseline (Table II(a), row 5), all using 128 GPUs. Adaptive compression with a 32 $\times$  sequence length reduction yields up to a 7.1 $\times$  speedup with no loss in PSNR or SSIM. Further compression yields diminishing returns due to quad-tree overhead. Tiling provides up to 1.9 $\times$  with 16 tiles per sample. Further tiling introduces excessive halo padding overhead and degrades computing performance. Accuracy remains stable across all settings.

### B. Maximal Sequence Length Scaling

Table III presents sequence length and resolution scaling performance of various model architectures and strategies, demonstrating how the combination of spatial compression, tiling, and the Reslim architecture enables extreme sequence lengths. We achieve sequence lengths of up to 4.2 billion tokens (global downscaling resolution of 0.9 km) for a 9.5M parameter model and up to 671 million tokens (global resolution 2.3 km) for a 10B parameter model. These results surpass the state-of-the-art in sequence scaling by more than 22,000 $\times$ , compared to state-of-the-art sequence parallelism of 188K tokens [22], and the Swin Transformer at 147K tokens [27].

All experiments utilize 23 input variables (12 atmospheric, 6 surface, and 5 static) and produce 18 output variables (excluding static inputs). Using a standard ViT with 9.5M parameters, the maximum sequence length is limited to 25K tokens (coarse 156 km global resolution) when using 8 GPUs. Scaling this ViT model to 10B parameters results in an out-of-memory (OOM) error, making global downscaling infeasible.

In contrast, Reslim demonstrates significantly better scaling. With just 8 GPUs, a 9.5M parameter Reslim model scales to 298M tokens at a 3.5 km global resolution. This corresponds to an output tensor of shape [5760, 4520, 18], assuming a 2 $\times$ 2 image patch size. Increasing the number of GPUs to 32, we achieve 466M tokens at 2.7 km resolution.


When combining Reslim with both spatial tiling (16 tiles per sample) and adaptive spatial compression (4 $\times$ ) techniques, substantial improvements are obtained. With these methods, the model achieves 1.1B tokens on only 8 GPUs, corresponding to downsampled output of size [11520, 23040, 18]. This result is made possible through several key compression techniques:

- Channel aggregation in Reslim (see Fig. 2) reduces the sequence length by 18 $\times$  by aggregating channels.
- Spatial tiling divides the sample into 16 tiles, reducing the sequence length per GPU by 16 $\times$ .
- Adaptive spatial compression reduces sequence by 4 $\times$ .
- Reslim processes directly on low-resolution input, reducing the effective sequence length by 60 $\times$ .

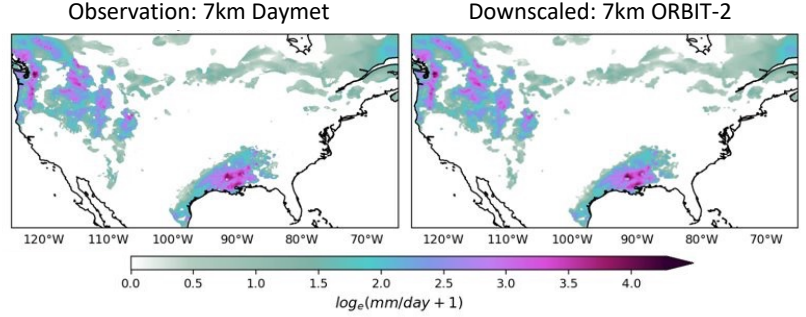
Combining all four, the effective per-GPU sequence length becomes only 17,280 tokens, despite the global output representing 1.1 billion tokens. Finally, by scaling to 128 GPUs, we achieve our largest configuration: 4.2 billion tokens at 0.9 km resolution.

For the 10B parameter model, Reslim still scales efficiently. Without compression or tiling, it reaches 18 million tokens. With 4 $\times$  compression, 16 tiles, and 512 GPUs, the model handles 671 million tokens at 2.3 km resolution.





(a) Power spectrum of minimum temperature downscaling



(b) Daily total precipitation vs. observation

Fig. 7: (a) Power spectra of downsampled minimum temperature using 9.5M and 126M parameter models, showing improved high-frequency fidelity with increased model capacity. (b) Daily total precipitation over the U.S. on January 1, 2020: (Left) 7 km DAYMET ground truth, and (Right) 7 km ORBIT-2 downscaling using the 126M model. ORBIT-2 accurately reconstructs fine-scale precipitation structures.

### C. TILES Sequence Scaling Speedup

Fig. 6(a) demonstrates the scalability of the TILES algorithm. With 16 tiles per sample, TILES achieves a 1.9 $\times$  speedup over the non-tiling baseline, both at 8 GPUs, using a 9.5M parameter model on the ERA5 $\rightarrow$ ERA5 112 $\rightarrow$ 28 km downscaling task. As GPU count increases, speedup scales nearly linearly, reaching 515 $\times$  at 2048 GPUs relative to the 8-GPU baseline without tiling. This highlights the scalability and minimal overhead of the TILES approach for distributed training.

### D. Strong Scaling Efficiencies & Throughput

Fig. 6(b) presents strong scaling performance across model sizes (9.5M to 10B parameters) using the same dataset as in Fig. 6(a). Experiments were conducted at scales of 64, 256, 1024 and 4096 nodes with 8 GPUs for each node, employing all four forms of orthogonal parallelism (see Sec. III). Each point in the figure reports the average runtime in second per hourly sample with a data label for corresponding strong scaling efficiency, relative to each model’s baseline performance at 64 nodes (512 GPUs).

All model sizes maintain high strong scaling efficiencies between 92–98%. The smallest model (9.5M) underutilizes hardware at large scales due to insufficient computing with small model, with 2.5e-6 seconds per sample and a sustained computing throughput at 363 PetaFLOPS at 4096 nodes (32,768 GPUs). In contrast, larger models saturate compute resources: the 126M, 1B, and 10B models reach sustained throughputs of 1.3, 1.5, and 1.8 ExaFLOPS, respectively, at 4096 nodes. These results demonstrate the strong scalability of Reslim and the effectiveness of orthogonal parallelism for exascale climate downscaling.

### E. Fine-Tuning & Inference

US Regional Fine-Tuning. Following pretraining on the datasets in Table I, we fine-tune ORBIT-2 models with 9.5M and 126M parameters on two tasks: (1) US-specific downscaling of ERA5 and DAYMET data from 28 km to 7

km, evaluated against 7 km DAYMET observations for both daily total precipitation and minimum temperature; and (2) global downscaling of 28 km ERA5 precipitation to 7 km, evaluated against IMERG observations. These tasks evaluate both regional accuracy and global generalization. The current fine-tuning results are limited to 9.5M and 126M models. We will update the manuscript soon to also report fine-tuning results for the 1B and 10B parameter models.

Table IV(a) summarizes results for daily minimum temperature downscaling. Both 9.5M and 126M models accurately reconstruct high-resolution temperature at 7 km resolution and are capable of capturing extremes, with the larger 126M model consistently outperforming the smaller 9.5M model across all metrics. Notably, the 126M model achieves an  $R^2$  of 0.999 and SSIM of 0.987, establishing a new benchmark for temperature downscaling at 7 km resolution.

Fig. 7(a) provides corresponding spectral analysis for the downscaling results in Table IV(a) by comparing the spatial power spectra of the two model sizes. The 126M model accurately captures high-frequency content, closely matching the high frequency spectral characteristics of the DAYMET observation ground truth. In contrast, the 9.5M model deviates from the ground truth at high frequencies. This demonstrates the larger model’s ability to resolve fine-scale spatial variability, emphasizing the value of increased model capacity for high-fidelity climate downscaling.

Table IV(b) presents downscaling results for daily total precipitation—one of the most challenging variables due to its high spatial variability and localized extremes. ORBIT-2 demonstrates strong performance, closely matching observations when downscaling to 7 km resolution. Fig. 7(b) provides a visual comparison of 7 km DAYMET ground truth and ORBIT-2’s 7 km prediction. Since ERA5 (a reanalysis product) and DAYMET (observations) have different statistical properties, we combine them at 28 km resolution as the input for downscaling to 7 km. The 126M model consistently outperforms the 9.5M model across all evaluation metrics, achieving an  $R^2$  of 0.979 and an overall RMSE of 0.135

Fig. 8: Playable animation comparing daily total precipitation at global scale for July 2020. (Left) Coarse-resolution ERA5 28km input; (Middle) IMERG 7km ground truth; (Right) ORBIT-2 downsampled prediction at 7 km. The animation demonstrates ORBIT-2’s ability to recover fine-scale features from coarse inputs. For optimal playback, please view in Adobe Acrobat.

mm/day. Notably, the large model also accurately captures extreme precipitation events, with RMSE values of 0.365 mm/day at the 99.7th percentile and 0.525 mm/day at the 99.99th percentile. All RMSE values for precipitation are computed in log-transformed space using  $\log(x+1)$ , where  $x$  denotes daily precipitation in millimeters.

Global Model Inference. Fig. 8 shows an animation of global daily precipitation for July 2020, comparing ORBIT-2 predictions (downscaling from 28 km to 7 km) against 7 km IMERG observations. The model was applied directly for global inference—without any fine-tuning or bias correction—to evaluate generalization. Despite challenges in downscaling ERA5 reanalysis data to satellite IMERG data due to their data source inconsistency, ORBIT-2 demonstrates strong generalization capability, achieving  $R^2 = 0.90$ , SSIM = 0.96, PSNR = 41.8, and RMSE = 0.34 mm/day (in  $\log(x+1)$  space). These results underscore ORBIT-2’s robustness, effectively extending regional training to global inference.

## VI. IMPLICATION

ORBIT-2 represents a significant leap forward in the convergence of artificial intelligence, HPC, and Earth system science. By overcoming fundamental challenges in scalability, resolution, and uncertainty, ORBIT-2 sets a new standard for climate and earth system foundation models.

Impact on HPC. ORBIT-2 pushes the frontiers of HPC by enabling ViTs at unprecedented scale. Through its novel TILES algorithm, ORBIT-2 reduces the self-attention complexity from quadratic to linear, allowing for efficient processing of ultra-long sequences. Complementing this, the Reslim architecture introduces a lightweight, uncertainty-aware learning framework that leverages residual learning and Bayesian regularization to improve efficiency and training robustness.

Crucially, ORBIT-2 achieves breakthrough scalability in ViT training, scaling up to 10 billion parameters model size across 32,768 GPUs, and scale up to 4.2 billion token

sequence length, several magnitudes longer than the state-of-the-art long sequence implementation that scales to 188K tokens. It sustains up to 1.8 ExaFLOPs of performance with 92–98% strong scaling efficiency at 32,768 GPUs, setting a new benchmark for exascale AI workloads. ORBIT-2 serves as a blueprint for next-generation exascale foundation models, enabling transformative applications across domains including genomics, fluid dynamics, astrophysics, and Earth system modeling.

Impact on Earth System Science. For Earth system modeling, ORBIT-2 enables hyper-resolution, global-scale downscaling with state-of-the-art accuracy and efficiency. Evaluated against observational data, it achieves a coefficient of determination ( $R^2$ ) of 0.999 for temperature and 0.979 for precipitation at 7 km resolution across the continental United States, closely matching ground-truth observation records. Critically, ORBIT-2 generalizes across both variables and geographic regions using a single unified foundation model—an essential capability for improving climate projection fidelity, detecting localized extremes, and guiding mitigation and adaptation strategies. This makes ORBIT-2 a valuable tool for researchers, policy-makers, and emergency planners worldwide.

## ACKNOWLEDGMENTS

This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The U.S. Government retains a nonexclusive, worldwide license to publish or reproduce the published form of this manuscript, or to authorize others to do so, for U.S. Government purposes, as acknowledged by the publisher. This research was primarily supported by the ORNL’s AI Initiative sponsored by the Director’s Research and Development Program at ORNL, additionally supported by DOE Early Career Project sponsored by the BER program. An award of computer time was provided by the INCITE program. This research also used resources of the Oak Ridge Leadership

Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

## REFERENCES

- [1] F. Giorgi, "Thirty years of regional climate modeling: Where are we and where are we going next?" *Journal of Geophysical Research: Atmospheres*, vol. 124, no. 11, pp. 5696–5723, 2019.
- [2] P. G. Jones and P. K. Thornton, "Generating downscaled weather data from a suite of climate models for agricultural modelling applications," *Agricultural Systems*, vol. 114, pp. 1–5, 2013.
- [3] E. Gutmann, T. Pruitt, M. P. Clark, L. Brekke, J. R. Arnold, D. A. Raff, and R. M. Rasmussen, "An intercomparison of statistical downscaling methods used for water resource assessments in the united states," *Water Resources Research*, vol. 50, no. 9, pp. 7167–7186, 2014.
- [4] E. H. Girvetz, E. P. Maurer, P. B. Duffy, A. Ruesch, B. Thrasher, and C. Zganjar, "Making climate data relevant to decision making: the important details of spatial and temporal downscaling," *World Bank*, 2013.
- [5] I. González-Aparicio, F. Monforti, P. Volker, A. Zucker, F. Careri, T. Huld, and J. Badger, "Simulating european wind power generation applying statistical downscaling to reanalysis data," *Applied Energy*, vol. 199, pp. 155–168, 2017.
- [6] B. C. Hewitson and R. G. Crane, "Climate downscaling: techniques and application," *Climate Research*, vol. 7, no. 2, pp. 85–95, 1996.
- [7] F. Giorgi and W. J. Gutowski Jr, "Regional dynamical downscaling and the cordex initiative," *Annual review of environment and resources*, vol. 40, no. 1, pp. 467–490, 2015.
- [8] R. L. Wilby and C. W. Dawson, "The statistical downscaling model: insights from one decade of application," *International Journal of Climatology*, vol. 33, no. 7, 2013.
- [9] R. L. Wilby, T. Wigley, D. Conway, P. Jones, B. Hewitson, J. Main, and D. Wilks, "Statistical downscaling of general circulation model output: A comparison of methods," *Water resources research*, vol. 34, no. 11, pp. 2995–3008, 1998.
- [10] N. Koldunov, T. Rackow, C. Lessig, S. Danilov, S. K. Cheedela, D. Sidorenko, I. Sandu, and T. Jung, "Emerging ai-based weather prediction models as downscaling tools," *preprint arXiv:2406.17977*, 2024.
- [11] B. Kumar, R. Chattopadhyay, M. Singh, N. Chaudhari, K. Kodari, and A. Barve, "Deep learning-based downscaling of summer monsoon rainfall data over indian region," *Theoretical and Applied Climatology*, vol. 143, pp. 1145–1156, 2021.
- [12] F. Wang, D. Tian, L. Lowe, L. Kalin, and J. Lehrter, "Deep learning for daily precipitation and temperature downscaling," *Water Resources Research*, vol. 57, no. 4, p. e2020WR029308, 2021.
- [13] N. J. Annau, A. J. Cannon, and A. H. Monahan, "Algorithmic hallucinations of near-surface winds: Statistical downscaling with generative adversarial networks to convection-permitting scales," *Artificial Intelligence for the Earth Systems*, vol. 2, no. 4, p. e230015, 2023.
- [14] V. Eyring, P. Gentile, G. Camps-Valls, D. M. Lawrence, and M. Reichstein, "Ai-empowered next-generation multiscale climate modelling for mitigation and adaptation," *Nature Geoscience*, vol. 17, no. 10, pp. 963–971, 2024.
- [15] Y. Liu, J. Doss-Gollin, G. Balakrishnan, and A. Veeraraghavan, "Generative precipitation downscaling using score-based diffusion with wasserstein regularization," *arXiv preprint arXiv:2410.00381*, 2024.
- [16] M. Aich, P. Hess, B. Pan, S. Bathiany, Y. Huang, and N. Boers, "Conditional diffusion models for downscaling & bias correction of earth system model precipitation," *arXiv preprint arXiv:2404.14416*, 2024.
- [17] X. Wang, S. Liu, A. Tsaris, J.-Y. Choi, A. Aji, M. Fan, W. Zhang, J. Yin, M. Ashfaq, D. Lu, and P. Balaprakash, "Orbit: Oak ridge base foundation model for earth system predictability," 2024. [Online]. Available: <https://arxiv.org/abs/2404.14712>
- [18] J. Schmude *et al.*, "Prithvi wx: Foundation model for weather and climate," 2024. [Online]. Available: <https://arxiv.org/abs/2409.13598>
- [19] T. Nguyen, J. Jewik, H. Bansal, P. Sharma, and A. Grover, "Climatelearn: Benchmarking machine learning for weather and climate modeling," 2023.
- [20] D. Han, X. Pan, Y. Han, S. Song, and G. Huang, "Flatten transformer: Vision transformer using focused linear attention," 2023. [Online]. Available: <https://arxiv.org/abs/2308.00442>
- [21] S. A. Jacobs, M. Tanaka, C. Zhang, M. Zhang, S. L. Song, S. Rajbhandari, and Y. He, "DeepSpeed Ulysses: System optimizations for enabling training of extreme long sequence transformer models," 2023. [Online]. Available: <https://arxiv.org/abs/2309.14509>
- [22] A. Tsaris, C. Zhang, X. Wang, J. Yin, S. Liu, M. Ashfaq, M. Fan, J. Y. Choi, M. Wahib, D. Lu, P. Balaprakash, and F. Wang, "Sequence length scaling in vision transformers for scientific images on frontier," 2024. [Online]. Available: <https://arxiv.org/abs/2405.15780>
- [23] H. Li, Y. Yang, M. Chang, S. Chen, H. Feng, Z. Xu, Q. Li, and Y. Chen, "Srdiff: Single image super-resolution with diffusion probabilistic models," *Neurocomputing*, vol. 479, pp. 47–59, 2022.
- [24] Z. Lu, J. Li, H. Liu, C. Huang, L. Zhang, and T. Zeng, "Transformer for single image super-resolution," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 457–466.
- [25] T. Nguyen, J. Brandstetter, A. Kapoor, J. K. Gupta, and A. Grover, "Climax: A foundation model for weather and climate," *arXiv preprint arXiv:2301.10343*, 2023.
- [26] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," 2021. [Online]. Available: <https://arxiv.org/abs/2103.14030>
- [27] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, F. Wei, and B. Guo, "Swin transformer v2: Scaling up capacity and resolution," 2022. [Online]. Available: <https://arxiv.org/abs/2111.09883>
- [28] Z. Tu, H. Talebi, H. Zhang, F. Yang, P. Milanfar, A. Bovik, and Y. Li, "Maxvit: Multi-axis vision transformer," 2022. [Online]. Available: <https://arxiv.org/abs/2204.01697>
- [29] X. Wang, I. Lyngaas, A. Tsaris, P. Chen, S. Dash, M. C. Shekar, T. Luo, H.-J. Yoon, M. Wahib, and J. Gouley, "Ultra-long sequence distributed transformer," 2023. [Online]. Available: <https://arxiv.org/abs/2311.02382>
- [30] I. Lyngaas, M. G. Meena, E. Calabrese, M. Wahib, P. Chen, J. Igarashi, Y. Huo, and X. Wang, "Efficient distributed sequence parallelism for transformer-based image segmentation," *Electronic Imaging*, vol. 36, no. 12, pp. 199–1–199–1, 2024.
- [31] Y. Zhao *et al.*, "Pytorch fsdp: Experiences on scaling fully sharded data parallel," *Proceeding of VLDB Endowment*, vol. 16, no. 12, p. 3848–3860, August 2023.
- [32] M. Shoyebi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-lm: Training multi-billion parameter language models using model parallelism," 2020.
- [33] C. He, S. Li, M. Soltanolkotabi, and S. Avestimehr, "Pipetransformer: Automated elastic pipelining for distributed training of transformers," 2021.
- [34] Y. Huang, Y. Cheng, A. Bapna, O. Firat, M. X. Chen, D. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, and Z. Chen, "Gpipe: Efficient training of giant neural networks using pipeline parallelism," 2019.
- [35] C. Kim, H. Lee, M. Jeong, W. Baek, B. Yoon, I. Kim, S. Lim, and S. Kim, "torchgpipe: On-the-fly pipeline parallelism for training giant models," 2020.
- [36] E. Zhang, I. Lyngaas, P. Chen, X. Wang, J. Igarashi, Y. Huo, M. Wahib, and M. Munetomo, "Adaptive patching for high-resolution image segmentation with transformers," 2024. [Online]. Available: <https://arxiv.org/abs/2404.09707>
- [37] C. Huang, J. R. Townshend, S. Liang, S. N. Kalluri, and R. S. DeFries, "Impact of sensor's point spread function on land cover characterization: assessment and deconvolution," *Remote Sensing of Environment*, vol. 80, no. 2, pp. 203–212, 2002.
- [38] Q. Wang, Y. Tang, and P. M. Atkinson, "The effect of the point spread function on downscaling continua," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 168, pp. 251–267, 2020.
- [39] Q. Wang, C. Zhang, X. Tong, and P. M. Atkinson, "General solution to reduce the point spread function effect in subpixel mapping," *Remote Sensing of Environment*, vol. 251, p. 112054, 2020.
- [40] S. Li *et al.*, "Pytorch distributed: Experiences on accelerating data parallel training," *Proceedings of the VLDB Endowment*, vol. 13, 2020.
- [41] T. Dao, "Flashattention-2: Faster attention with better parallelism and work partitioning," 2023.
- [42] Pytorch, "Automatic mixed precision package - torch.amp," <https://pytorch.org/docs/stable/amp.html#torch.cuda.amp.GradScaler>, 2024, accessed: 2024-04-05.
- [43] H. Hersbach *et al.*, "The era5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020.
- [44] C. Daly and K. Bryant, "The prism climate and weather system—an introduction," *Corvallis, OR: PRISM climate group*, vol. 2, 2013.
- [45] P. E. Thornton, M. M. Thornton, B. W. Mayer, N. Wilhelmi, Y. Wei, R. Devarakonda, and R. B. Cook, "Daymet: Daily surface weather data on a 1-km grid for north america, version 2," Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), Tech. Rep., 2014.
- [46] G. J. Huffman *et al.*, "Integrated multi-satellite retrievals for the global precipitation measurement (gpm) mission (imerg)," *Satellite precipitation measurement: Volume 1*, pp. 343–353, 2020.
- [47] Microsoft, "Microsoft DeepSpeed," <https://github.com/microsoft/DeepSpeed>, 2024, accessed: 2024-04-05.