

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TRANSFORMING WEATHER DATA FROM PIXEL TO LATENT SPACE

Anonymous authors

Paper under double-blind review

## ABSTRACT

The increasing impact of climate change and extreme weather events has spurred growing interest in deep learning for weather research. However, existing studies often rely on weather data in pixel space, which presents several challenges such as smooth outputs in model outputs, limited applicability to a single pressure-variable subset (PVS), and high data storage and computational costs. To address these challenges, we propose a novel Weather Latent Autoencoder (WLA) that transforms weather data from pixel space to latent space, enabling efficient weather task modeling. By decoupling weather reconstruction from downstream tasks, WLA improves the accuracy and sharpness of weather task model results. The incorporated Pressure-Variable Unified Module transforms multiple PVS into a unified representation, enhancing the adaptability of the model in multiple weather scenarios. Furthermore, weather tasks can be performed in a low-storage latent space of WLA rather than a high-storage pixel space, thus significantly reducing data storage and computational costs. Through extensive experimentation, we demonstrate its superior compression and reconstruction performance, enabling the creation of the ERA5-Latent dataset with unified representations of multiple PVS from ERA5 data. The compressed full PVS in the ERA5-Latent dataset reduces the original 244.34 TB of data to 0.43 TB. The downstream task further demonstrates that task models can apply to multiple PVS with low data costs in latent space and achieve superior performance compared to models in pixel space.

## 1 INTRODUCTION

The profound impact of climate change and extreme weather events on the Earth has attracted widespread attention (Patz et al., 2005; Wild et al., 2025; Chen et al., 2025). Recently, deep learning methods have made groundbreaking advancements in meteorology, leading to increasing interest in their application to weather research (Ravuri et al., 2021; LIU et al., 2022; Yang et al., 2023; Zhang et al., 2023b; Gong et al., 2024a;b). However, most existing studies focus primarily on weather-related tasks in the pixel space of weather data (Bi et al., 2023; Chen et al., 2023b;a). The efficiency of weather models in prior studies is often hindered by the inherent uncertainty of tasks and the diversity of data, whereas data costs are inflated by expensive storage and processing requirements.

Specifically, performing weather-related tasks in the pixel space presents three main limitations (as shown in Fig.1): 1) **Smooth Model Results**. Weather data contain rich small-scale extreme values. When performing tasks such as weather forecasting and downscaling in the pixel space, the model also needs to perform weather reconstruction, requiring a fine reconstruction of small-scale extreme values. However, the inherent uncertainty in weather-related tasks degrades the performance of small-scale extreme values reconstruction and extreme events prediction, leading to smooth results (Ravuri et al., 2021). 2) **Limited Model Applicability to a Single Pressure-Variable Subset (PVS)**. Weather data typically record various weather variables across multiple pressure levels, leading to significant data diversity in the pixel space (Astruc et al., 2024; Xiong et al., 2024). Different weather-related tasks and applications often require distinct PVS selections. For instance, the 500 hPa geopotential height and the 850 hPa wind fields are fundamental in representing atmospheric steering flows and vortex dynamics, which are key to typhoon path prediction (Hua & Chong-Yin, 2010; Moore & Dixon, 2015). Conversely, the 500 hPa geopotential height, 700 hPa vertical velocity, and 925 hPa specific humidity serve as essential parameters for short-term rainfall forecasting (Kuligowski & Barros, 1998; Tian et al., 2015). However, models trained in pixel space are typically restricted to a single PVS,

limiting their adaptability across multiple weather scenarios requiring different PVS compositions. 3) **High Data Storage and Computational Costs.** Pixel-based weather datasets can reach hundreds of terabytes (TB) or even petabytes (PB), leading to significant storage and computational costs (Hersbach et al., 2020). This poses a substantial challenge for the large-scale application of deep learning in meteorology (Klöwer et al., 2021; Han et al., 2024a).

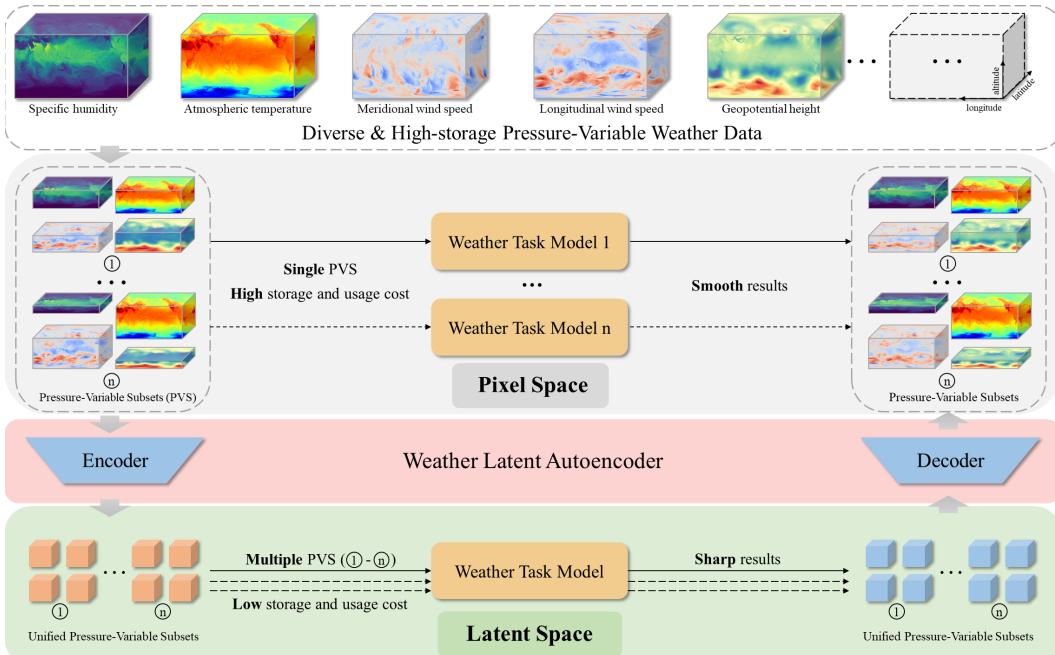


Figure 1: Transforming weather data from diverse and high-storage pixel space to unified and low-storage latent space for weather tasks using weather latent autoencoder. The weather task model in pixel space suffers from high data storage and computational costs and limited applicability to single pressure-variable subset, often yielding ambiguous results. In contrast, the model in latent space benefits from reduced data storage and computational costs, enabling the use of multiple pressure-variable subsets and producing sharper results.

To address the above limitations, we propose a novel approach that transforms weather data from pixel space to latent space for weather-related tasks. Specifically, we introduce the Weather Latent Autoencoder (WLA), as illustrated in Fig.1. WLA effectively encodes diverse and high-storage weather data from the pixel space to a unified and lower-storage latent space, facilitating its application to multiple PVSs. This transformation allows weather-task models to operate directly in the latent space, eliminating the need for pixel-space data, thereby enhancing their adaptability to different PVS compositions while significantly reducing data storage and computational costs.

Specifically, WLA addresses the aforementioned issues in three ways: 1) **Decoupling Weather Reconstruction from Weather Tasks**. In this approach, weather tasks are performed in the latent space, while weather reconstruction occurs within the pretrained Weather Latent Autoencoder. The pretrained WLA ensures that latent features effectively preserve small-scale extreme values, allowing for high-quality reconstruction from these features. During weather tasks in the latent space, the uncertainty inherent in these tasks has minimal impact on the small-scale extreme values reconstruction of WLA, resulting in sharp and accurate outcomes for the weather task model. 2) **Unified Pressure-Variable Representation**. We introduce a Pressure-Variable Unified Module (PVUM) designed to transform any pressure-variable subset to a unified space. PVUM leverages pressure-variable metadata in weather data to generate adaptive mapping network weights through a hypernetwork, enabling the conversion of weather data from pixel space into a unified latent space. This framework allows the weather task model to seamlessly accommodate various types of weather data inputs in the latent space, enhancing its applicability across diverse weather scenarios. 3) **Latent Space Framework**. We propose the Latent Space Framework, which transitions weather task models

108 from pixel space to latent space, significantly reducing data storage and computational costs. Thanks  
 109 to WLA’s superior compression and reconstruction capabilities, the latent data retain most of the  
 110 information from the original pixel data, but with a much smaller storage footprint. This results in  
 111 a substantial reduction in storage costs. Furthermore, tasks such as model training, validation, and  
 112 testing, which typically require large amounts of data, can be carried out using low-storage latent  
 113 data, yielding significant savings in data computational costs.

114 To facilitate research on weather tasks performed directly in latent space, we introduce **ERA5-Latent**,  
 115 a novel dataset derived from ERA5 (Hersbach et al., 2020). While raw ERA5 data offers high fidelity,  
 116 its sheer size (hundreds of TB) is prohibitive. Common alternatives like Weatherbench provide  
 117 fixed, lower-resolution (128×256 size) subsets with limited variables, restricting their suitability for  
 118 contemporary deep learning applications. Addressing these challenges, we utilize our proposed WLA  
 119 to transform high-resolution ERA5 data (721×1440 size) from pixel space into a compact latent  
 120 representation. This ERA5-Latent dataset substantially reduces data costs and enables research using  
 121 the full scope of ERA5 maps and diverse variable sets within the latent domain.

122 The original ERA5 data includes 164 variables and totals 244.34 TB. Our WLA transformation  
 123 compresses this into a latent representation requiring only 0.43 TB, achieving a 566.3× compression  
 124 ratio and significantly lowering storage costs. To support diverse modeling needs, ERA5-Latent  
 125 offers unified latent representations for commonly used configurations: 6 upper-air variables across 6,  
 126 13, and 25 pressure levels; surface variables in sets of 4 and 8; and precipitation variables in sets of 1  
 127 and 6. Models can leverage this low-storage latent data for training, validation, and testing across  
 128 various scenarios in latent space, minimizing data and computational expenses.

129 In summary, our main contributions are as follows:

- 131 1. We propose a novel framework that transforms weather data from pixel space to latent  
 132 space for weather tasks. By transforming data into latent space, we decouple weather  
 133 reconstruction from the downstream tasks, enabling the model to generate sharp and accurate  
 134 results. The unified representation in the latent space allows task models to handle multiple  
 135 pressure-variable subsets, while the latent representation significantly reduces data storage  
 136 and computational costs.
- 137 2. We introduce the Weather Latent Autoencoder for the pixel-to-latent transformation of  
 138 weather data. WLA can effectively transform any pressure-variable subset from pixel space  
 139 to a unified latent space, providing excellent compression and reconstruction performance.  
 140 It can be combined with weather forecasting models that integrate advanced techniques,  
 141 promoting weather research from a data optimization perspective.
- 142 3. We have constructed the ERA5-Latent dataset, which provides large-scale ERA5 weather  
 143 data with multiple pressure-variable subsets in a smaller data storage footprint and unified  
 144 latent space.

## 146 2 RELATED WORK

### 148 2.1 WEATHER DATA COMPRESSION

149 Weather data compression has advanced from traditional linear quantization (GRIB2-based 17×  
 150 compression of CAM (Inness et al., 2019; Klöwer et al., 2021)) to neural representation learning.  
 151 Autoencoder-based models (Liang et al., 2023) and coordinate-aware networks (Huang & Hoefler,  
 152 2023) achieve high compression ratios through instance-specific overfitting, though often at the  
 153 cost of generalization. Meta-learning methods like COIN++ (Dupont et al., 2022) address this  
 154 by leveraging shared priors for modality-agnostic compression. More recent advances combine  
 155 probabilistic modeling with entropy coding; for example, Mirowski et al. (Mirowski et al., 2024)  
 156 achieve 1000× compression using hyperpriors and vector quantization, while CRA5 (Han et al.,  
 157 2024a) employs a dual-variational transformer to optimize rate-distortion via hierarchical latent space  
 158 modeling. These compression methods have limitations when using meteorological compressed  
 159 data: either online decompression of data leads to computational resource consumption, or offline  
 160 decompression leads to storage resource consumption. Our weather latent autoencoder avoids the  
 161 data decompression step by directly using latent space data, reducing data calculation costs.

162  
163

## MANAGING THE DIVERSITY OF WEATHER DATA

164  
165  
166  
167  
168  
169  
170  
171  
172  
173

Earth science modeling is challenged by heterogeneous observational data. Current methods either rely on specialized architectures such as Omnisat’s modality-specific encoders for cross-modal feature alignment (Astruc et al., 2024) or on metadata-driven adaptation, as seen in DOFA’s spectral self-supervision (Xiong et al., 2024). In weather forecasting, the combinatorial complexity of atmospheric variables and pressure levels often results in brittle models. For instance, FengWu (Chen et al., 2023a) employs 5 upper-air variables at 37 pressure levels with 4 surface variables, while Pangu (Bi et al., 2023) and FengWu-GHR (Han et al., 2024b) use 13 pressure levels for similar variables. FuXi (Chen et al., 2023b) uses 5 upper-air variables (13 levels) with an expanded set of 5 surface variables, and Gencast (Price et al., 2024) scales to 6 variables each. These differences underscore the need for unified frameworks that can flexibly handle diverse pressure-variable subsets.

174  
175

## LOW-COST WEATHER DATASETS

176  
177  
178  
179  
180  
181  
182  
183

The exponential growth of weather data poses significant challenges in storage, computation, and accessibility. Curated low-cost datasets such as Weatherbench (Rasp et al., 2020) mitigate these issues by downsampling ERA5 reanalysis data to a  $1.405^\circ$  resolution ( $128 \times 256$ ) with 13 pressure levels, cutting storage requirements by 94% compared to native resolutions. In contrast, CRA5 retains ERA5’s full  $0.25^\circ$  resolution ( $721 \times 1440$ ) across 159 fields, achieving similar storage efficiency through neural compression at the expense of requiring decoder reconstruction. Collectively, these studies highlight the importance of efficient data representation and unified frameworks for advancing both atmospheric modeling and computer vision applications.

184  
185

## 3 METHOD

186

## 3.1 OVERVIEW OF WEATHER LATENT AUTOENCODER

187  
188  
189  
190  
191  
192  
193  
194

The Weather Latent Autoencoder transforms weather data from diverse and high-storage pixel space into unified and low-storage latent space. As illustrated in Figure 2, our framework integrates three core components: (1) a Pressure-Variable Unified Module that leverages metadata information to align heterogeneous PVS features, (2) a VAEformer Encoder-Decoder pair adopting the transformer architecture from CRA5’s pretraining stage (Han et al., 2024a) for latent feature compression/reconstruction, and (3) a Binary Quantization Module (BQM) that generates compact bitwise tokens through spherical normalization and binary quantization.

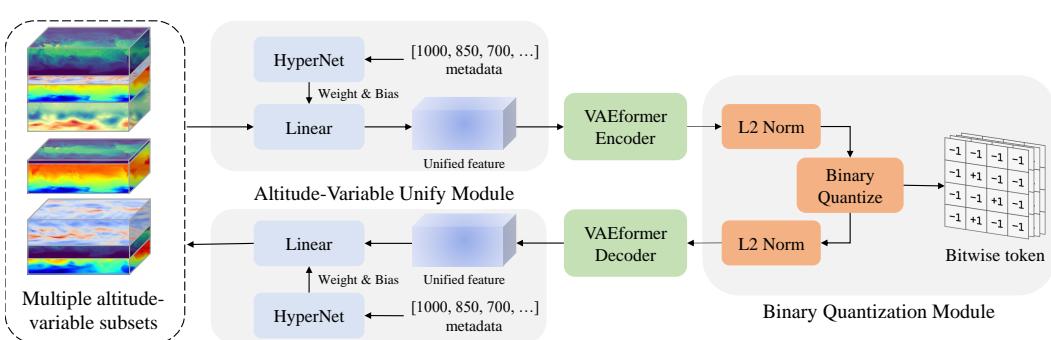
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206207  
208  
209

Figure 2: Architecture of the Weather Latent Autoencoder, which compresses weather data from a diverse, high-storage pixel space into a unified, low-storage latent space, and reconstructs it back into the pixel space.

210  
211  
212  
213  
214  
215

During the compression phase, the model starts with selecting multiple PVSs from a multiple pressure-variables weather dataset. The PVUM first converts pressure-variable metadata into adaptive parameters through a hypernetwork, enabling cross-scale feature alignment across disparate PVS. These unified features are subsequently encoded by the VAEformer encoder into low-dimensional latent representations, preserving essential weather patterns while discarding pixel-space redundancies. The BQM then projects the latent features onto a unit spherical space through L2-normalization and

216 applies binary quantization to produce storage-efficient bitwise tokens. This compression effectively  
 217 reduces data storage compared to original pixel-space PVS representations.  
 218

219 The reconstruction phase executes an inverse transformation through three cascaded operations.  
 220 Initially, bitwise tokens are mapped back to spherical space via L2-normalization. The VAEformer  
 221 Decoder subsequently reconstructs unified features through upsampling operators, ensuring high-  
 222 quality weather data reconstruction. Finally, the PVUM regenerates the original PVS by applying  
 223 metadata-guided inverse transformations, thereby completing the latent-to-pixel space transformation  
 224 cycle.

225 The latent space framework of WLA offers three fundamental benefits. First, the unified encoding  
 226 enables weather-task models to directly operate on a unified latent space, eliminating structural  
 227 modifications for cross-PVS generalization. Second, the WLA decouples weather reconstruction  
 228 from task modeling. The uncertainty present in the weather tasks does not affect the weather recon-  
 229 struction, ensuring that the task model can output sharp and accurate results. Third, data storage and  
 230 computational costs are significantly reduced as model training, validation, and inference primarily  
 231 utilize low-storage latent features, restricting pixel-space operations to final metric evaluation phases.  
 232 This layered approach addresses critical challenges in weather data processing, including multiple  
 233 pressure-variable representation learning, storage scalability, and task-specific adaptation.

### 236 3.2 PRESSURE-VARIABLE UNIFIED MODULE

238 To map any pressure-variable  
 239 subset from pixel space to a unified  
 240 feature, we designed the  
 241 Pressure-variable Unified Mod-  
 242 ule, which utilizes the metadata  
 243 of the pressure-variable subset to  
 244 generate adaptive weights and bi-  
 245 ases for a linear layer, thereby en-  
 246 abling adaptive feature mapping.

247 As shown in Figure 3, given  
 248 an input PVS tensor  $X \in \mathbb{R}^{C_1 \times H \times W}$  with its pressure-  
 249 variable metadata  $M \in \mathbb{R}^{C_1}$   
 250 (where  $C_1$  varies across tasks  
 251 and scenarios), PVUM generates a unified feature  $Y$  with fixed dimensionality through hypernetwork-  
 252 based parameter generation. This process consists of three core operations: **Metadata Embedding**:  
 253 The variable metadata  $M$  containing physical attributes (pressure levels and variables) undergoes  
 254 positional encoding followed by tokenization. A learnable class token [CLS] is prepended to the  
 255 token sequence  $T \in \mathbb{R}^{(C_1+1) \times d}$ , where  $d$  is the embedding dimension. **Cross-Variable Relation**  
 256 **Modeling**: The token sequence passes through several transformer blocks for learning the rela-  
 257 tionships between the metadata. **Adaptive Parameter Generation**: The [CLS] token produces bias  
 258 parameters  $b \in \mathbb{R}^{C_2}$  via a linear projection, while the remaining tokens generate a weight matrix  
 259  $W \in \mathbb{R}^{C_1 \times C_2}$  through another linear layer. The resulting weights  $W$  and bias  $b$  form a linear layer  
 260 that maps the features with  $C_1$  channels to features with  $C_2$  channels. Therefore, the input  $X$  is  
 261 reshaped from  $(C_1, H, W)$  to  $(L, C_1)$ , where  $L = H \times W$ , and then mapped to the target feature  $Y$   
 262 with shape  $(L, C_2)$  using the generated linear layer.

263 The PVUM structure not only adaptively maps any pressure-variable subset to a unified feature in  
 264 terms of shape, but also effectively preserves the relationships between weather data in pixel space.  
 265 Due to the continuity, smoothness, and vertical mixing of the atmosphere, there is inherent similarity  
 266 between different weather pressure levels and variable data, especially between adjacent pressure  
 267 levels for the same variable (Zhang et al., 2023a). The hypernetwork of PVUM learns this relationship  
 268 when modeling the metadata, allowing it to map similar weather variables and adjacent pressure  
 269 levels to similar unified features. As a result, PVUM preserves the relationships between weather  
 270 data in the feature space.

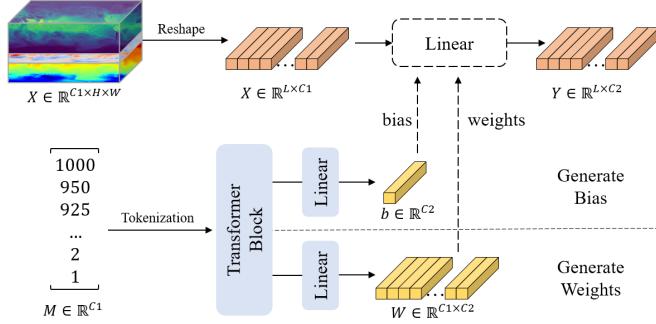


Figure 3: Workflow of Pressure-Variable Unified Module, which  
 transforms diverse weather data into unified representation.

270 3.3 BINARY QUANTIZATION MODULE  
271

272 To effectively compress weather features while preserving critical information, we propose the Binary  
273 Quantization Module that establishes a bi-directional mapping between continuous features and  
274 discrete binary tokens. As shown in Figure 2 (right), the module inherits the vector quantization  
275 framework from BSQ (Zhao et al., 2024) which has two key components: (1) spherical space  
276 projection for stable entropy loss estimation, and (2) deterministic binary quantization for hardware-  
277 friendly storage. The quantization process consists of three stages: First, input features undergo  
278 L2 normalization to project them onto a spherical space, which not only stabilizes the subsequent  
279 quantization but also enables computation of the entropy loss with acceptable memory/space cost  
280 (Han et al., 2025). Second, we apply sign-based binary quantization where positive values are mapped  
281 to 1 and negative values to -1, generating compact bitwise tokens. During reconstruction, the bitwise  
282 tokens are inversely projected to the spherical space through L2 normalization before being fed to the  
283 VAEformer decoder for upsampling.

284 The compression ratio of our weather latent autoencoder can be formally analyzed through the data  
285 storage. Let the input feature tensor  $F \in \mathbb{R}^{C \times H \times W}$  with float32 representation be compressed into bi-  
286 nary tokens  $B \in \{-1, 1\}^{C' \times H' \times W'}$ . The spatial downsampling factors  $(P_h, P_w) = (H/H', W/W')$   
287 combined with channel dimension adjustment yield a compression ratio

$$289 R = \frac{C \cdot H \cdot W \cdot 32}{C' \cdot H' \cdot W'} = \frac{C}{C'} \cdot 32 \cdot P_h \cdot P_w. \quad (1)$$

293 3.4 LATENT SPACE FRAMEWORK  
294

295 To reduce the data storage and compu-  
296 tational costs when using deep learning  
297 models for weather tasks, we propose the  
298 Latent Space Framework as illustrated in  
299 Figure 4. Given that the Weather Latent  
300 Autoencoder has effective compres-  
301 sion and reconstruction of weather data,  
302 our framework leverages two key obser-  
303 vations: (1) Latent representations pre-  
304 serve essential information from pixel-  
305 space data, and (2) Data similarity rela-  
306 tionships remain consistent across both  
307 pixel and latent spaces. Consequently,  
308 processes that require large amounts of data,  
309 such as training, validation, and testing of  
310 weather models, can be conducted in the  
311 lower-storage latent space.

312 Specifically, when conducting weather  
313 tasks, the input bitwise token passes  
314 through the weather task model, and the  
315 output prediction is compared to the target to compute the binary cross-entropy loss. During the  
316 training phase, the loss is used for gradient backpropagation and parameter updates. During the  
317 validation and testing phases, model performance can be directly evaluated on the low-storage latent  
318 space. In processes that only contain a small amount of data, such as calculating pixel metrics for  
319 weather tasks, we use the weather latent decoder to decode and reconstruct the predictions back to  
the high-storage pixel data, from which the corresponding pixel metrics can be calculated.

320 By utilizing the Latent Space Framework, we can store large amounts of data in the low-storage latent  
321 space and small amounts of data in the high-storage pixel space, effectively reducing data storage  
322 costs. Additionally, weather task models can be trained, validated, and tested in the latent space,  
323 where large amounts of data are required, while pixel space can be used for processes like metrics  
calculation that need only a small amount of data, thereby lowering the data computational costs.

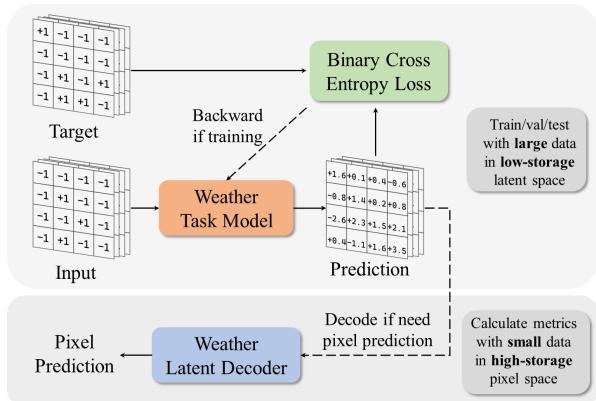


Figure 4: Overview of the Latent Space Framework. The data-intensive processes can be performed in the low-storage latent space, while processes requiring a smaller amount of data can be carried out in the high-storage pixel space, thereby effectively reducing data costs.

324 

## 4 EXPERIMENTAL RESULTS

325 

### 4.1 ORIGINAL DATASET

326 The ERA5 dataset (Hersbach et al., 2020) is serves as a standard for evaluating the Weather Latent  
 327 Autoencoder in comparison to other models. To meet various weather-related needs, we organized  
 328 three categories of variables: upper-air, surface, and precipitation variables. Considering the physical  
 329 characteristics of different weather variables, we treat them as different modalities and designed  
 330 distinct WLA architectures (Chen et al., 2023a): individual WLAs are trained for each upper-air  
 331 variable to capture multi-pressure-level dependencies, while unified WLAs handle surface and  
 332 precipitation variables to exploit intra-category correlations. Experimental details can be found in  
 333 Section A.1 of the supplementary materials.

334 

### 4.2 OVERALL RESULTS

335 To demonstrate the effectiveness of the WLA in data compression, we compared it with several  
 336 state-of-the-art compression methods (Elic (He et al., 2022), IEN (Xie et al., 2021), VQVAE (van den  
 337 Oord et al., 2017), VAGAN (Esser et al., 2021), VAEformer (Han et al., 2024a)) across three metrics:  
 338 compression ratio, bits per sub-pixel (bpsp) (Mentzer et al., 2019), and weighted RMSE (Han et al.,  
 339 2024b) on representative upper-air variables, surface variables, and precipitation variables. The  
 340 results are summarized in Table 1. Due to significant differences in numerical ranges among weather  
 341 variables, we included a "Variable Std" row as a reference. Generally, variables with higher variances  
 342 exhibit larger reconstruction errors. Since the number of input variables in WLA influences both  
 343 the compression ratio and reconstruction quality, we evaluated its performance using the maximum  
 344 input configurations: 25 pressure levels for upper-air variables, 8 variables for surface variables, and  
 345 6 variables for precipitation variables. The compression ratio and bpsp values reported for WLA in  
 346 Table 1 correspond to its performance on these three variable categories, respectively.

347 

350 Table 1: Compression Result of WLA and several state-of-the-art compression methods.

351 Method	352 Weighted RMSE ↓							353 Comp. Ratio ↑	354 bpsp ↓
	355 Upper-air Variables				356 Surface Variables		Precipitation		
	w500	w700	q700	q1000	TCC	SP	tp6h		
356 Var. Std (ref.)	0.218	0.240	0.0025	0.0059	0.36	9584.49	1.57	–	–
Elic (He et al., 2022)	0.197	0.233	0.00076	0.00087	0.18	537.82	1.19	648.3	0.112
IEN (Xie et al., 2021)	0.213	0.247	0.00084	0.00092	0.23	688.27	1.03	202.5	0.158
VQVAE (Mirowski et al., 2024)	0.382	0.401	0.00108	0.00113	0.19	673.32	1.29	<b>1100.0</b>	<b>0.029</b>
VQGAN (Mirowski et al., 2024)	0.367	0.371	0.00101	0.00107	0.18	652.38	1.20	<b>1100.0</b>	<b>0.029</b>
VAEformer (Han et al., 2024a)	0.117	0.134	0.00031	0.00035	0.12	376.90	0.80	323.1	0.099
<b>361 WLA (Upper-air)</b>	<b>0.076</b>	<b>0.083</b>	<b>0.00027</b>	<b>0.00028</b>	–	–	–	625.9	0.051
<b>WLA (Surface)</b>	–	–	–	–	<b>0.055</b>	<b>257.88</b>	–	200.3	0.159
<b>WLA (Precipitation)</b>	–	–	–	–	–	–	<b>0.47</b>	600.9	0.053

362 As shown in Table 1, WLA achieves superior overall compression performance across upper-air,  
 363 surface, and precipitation variables compared to existing methods, characterized by higher compression  
 364 ratios, lower bpsp values, and competitive weighted RMSE scores. This demonstrates that WLA  
 365 effectively balances weather data compression with reconstruction. Notably, WLA demonstrates  
 366 remarkable flexibility and versatility, enabling seamless adaptation to diverse variable combinations  
 367 and complex application scenarios. The visualization results and information loss analysis of WLA  
 368 can be found in Section A.5 of the supplementary materials.

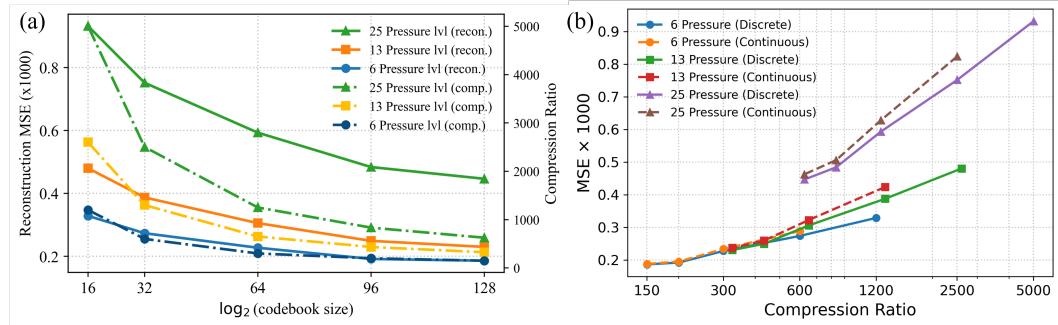
369 

### 4.3 OUT-OF-DOMAIN GENERALIZATION

370 To validate the generalization capability of WLA on unseen pressure levels and out-of-domain data,  
 371 we conducted experiments using two distinct datasets. The first is the ERA5 dataset, which includes  
 372 37 pressure levels, of which only 25 were used during the training phase. The second is the HRES  
 373 dataset, which was downsampled to match the spatial resolution of the ERA5 data.

378 The experimental results in Fig.5 show the performance of WLA on the temperature variable  
 379 across 12 pressure levels. These results indicate that WLA achieves robust reconstruction  
 380 performance on pressure levels for which it was not trained (denoted by \*). Furthermore,  
 381 the model shows a reconstruction performance on the HRES data that is comparable to its  
 382 performance on the ERA5 data. This demonstrates that WLA possesses strong generalization  
 383 capabilities, both for unseen pressure levels and for out-of-domain data. This indicates  
 384 that WLA can transform diverse weather data into a unified latent space representation, even for  
 385 out-of-domain data, thereby enabling the model in the latent space to naturally adapt to diverse data.  
 386  
 387  
 388  
 389  
 390  
 391

#### 392 4.4 ABLATION STUDY



405  
 406 Figure 6: **Ablation studies evaluated on the atmospheric temperature variable.** (a) Ablation study on  
 407 compression ratio and reconstruction quality of the WLA under varying input pressure levels (6, 13,  
 408 25 layers) and codebook sizes ( $2^{16}$  to  $2^{128}$ ). (b) Ablation study on the BQM, which can be seen as an  
 409 comparison of discrete and continuous latent space.

410 To identify the optimal balance between compression efficiency and reconstruction quality, we  
 411 conducted ablation studies on the atmospheric temperature variable using the upper-air dataset. We  
 412 evaluated the WLA across three input configurations (6, 13, and 25 pressure levels) and five codebook  
 413 sizes ( $2^{16}$ ,  $2^{32}$ ,  $2^{64}$ ,  $2^{96}$ ,  $2^{128}$ ).  
 414

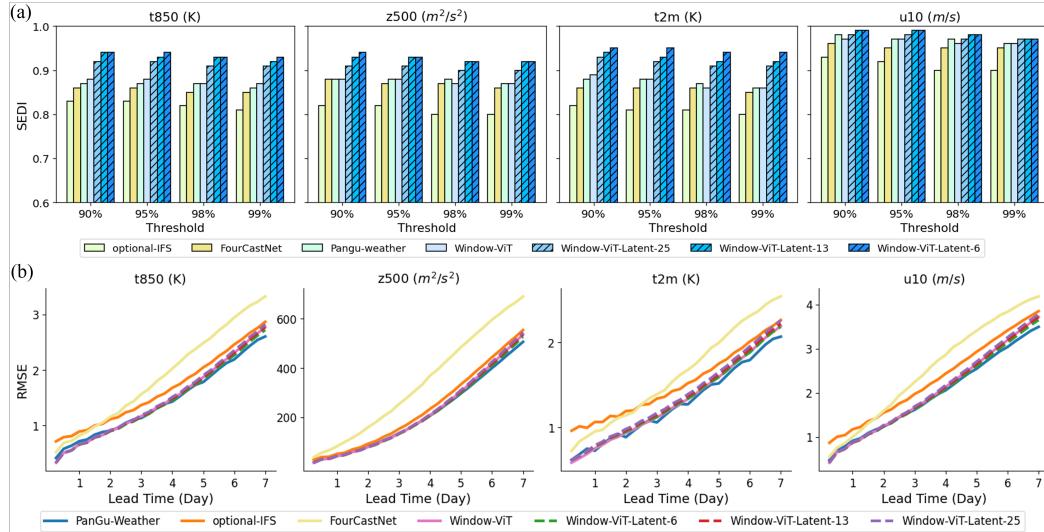
415 The general ablation results, illustrated in Fig. 6(a), demonstrate the inherent trade-off in our frame-  
 416 work: the compression ratio is inversely proportional to the codebook size but positively correlated  
 417 with the number of input pressure levels (consistent with Eq. 1). Conversely, reconstruction quality  
 418 improves with larger codebook sizes but decreases as the input data dimensionality grows. Based on  
 419 these observations, we selected a codebook size of  $2^{128}$  for our final model, as further increasing the  
 420 size yields diminishing returns in reconstruction quality.

421 To further investigate the role of discretization by BQM in data compression and reconstruction,  
 422 we compared our discrete framework against a continuous baseline. To ensure a fair comparison  
 423 under identical data compression rates, we constructed a continuous variant where the BQM is  
 424 replaced by a linear layer, which maps features to a latent space representation stored as a float32  
 425 with  $\log_2(\text{codebook size}) // 32$  channels. The channel dimensions were adjusted such that the total  
 426 bit-width remains constant. For example, for a codebook size of  $2^{128}$ , 128 binary channels are  
 427 compared against 4 float32 channels.

428 As shown in Fig.6(b), the discrete latent space achieves reconstruction performance comparable to the  
 429 continuous baseline at lower compression ratios. Notably, at higher compression ratios, the discrete  
 430 model significantly outperforms its continuous counterpart. This suggests that the BQM effectively  
 431 preserves essential semantic information during discretization, offering superior efficiency over a  
 432 bitrate-matched continuous representation.

432 4.5 DERIVED ERA5-LATENT DATASET  
433

434 Leveraging the excellent compression and reconstruction performance of the Weather Latent Au-  
435 toencoder, we transformed the multiple PVSs of ERA5 data into unified latent space, yielding the  
436 ERA5-Latent dataset. By utilizing the high compression rate of WLA, the ERA5-Latent dataset  
437 reduces the original 244.34 TB of data down to 0.43 TB, while providing a unified representation for  
438 multiple PVS. Dataset details can be found in subsection A.2.1 of supplementary materials.

439 440 4.6 DOWNSTREAM TASK VERIFICATION: WEATHER FORECASTING  
441

460 Figure 7: Comparison on the weather forecasting task. (a) Extreme weather forecasting performance  
461 on the SEDI metric. (b) RMSE for weather forecasting over a 7-day period. The number following  
462 "Window-ViT-latent" indicates the number of input pressure levels.

463  
464 To demonstrate that models operating in a unified latent  
465 space can adapt to multiple PVS and generate sharper  
466 results compared to pixel-space models, we conducted  
467 experiments on the weather forecasting task. We used  
468 two models to conduct weather forecasting: the Window-  
469 ViT (WT) in the pixel space, and the Window-ViT-  
470 Latent (WTL) in the latent space. The experiments were  
471 performed on the ERA5 dataset and the ERA5-Latent  
472 dataset. Specifically, following the settings in (Bi et al.,  
473 2023), we conduct 7-day weather forecasting at 6-hour  
474 intervals. For the WT, we use 13 pressure levels for five  
475 upper-air variables along with four surface variables. In  
476 contrast, to evaluate the WTL's adaptability to multiple  
477 PVS, we experiment with multiple pressure levels (25,  
478 13, 6) for the upper-air variables for the latent model.  
479 This experiment almost takes the same configuration as  
the state-of-the-art model (Chen et al., 2023a).

480 To comprehensively evaluate performance on extreme  
481 events, we employed two distinct metrics: the Sym-  
482 metricic Extremal Dependency Index (SEDI) (Kurth et al.,  
483 2023; Xu et al., 2024) to assess event detection capa-  
484 bility, and the Relative Quantile Error (RQE) (Kurth et al.,  
485 2023) to evaluate the preservation of extreme event magnitudes. We benchmarked our models against  
the physics-based Operational Integrated Forecasting System (IFS) (Bougeault et al., 2010) and the

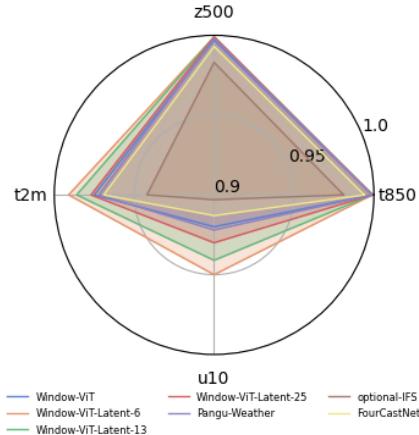


Figure 8: Extreme weather forecasting performance on the RQE metric. For easier visual comparison in the radar figure, the RQE is transformed using the formula  $1 + 10 \cdot \text{RQE}$ , with the optimal performance being closer to 1.

486 AI-based Pangu-Weather (Bi et al., 2023). Experimental details are provided in Section A.2 of the  
 487 supplementary materials.  
 488

489 **Superior preservation of extreme events.** As illustrated in Fig.7 (a) and Fig.8, WTL achieves  
 490 superior performance in extreme weather forecasting compared to the pixel-space baseline. The SEDI  
 491 scores confirm that WTL maintains high accuracy in detecting extreme occurrences. Furthermore,  
 492 the radar charts based on the RQE metric demonstrate that WTL effectively minimizes errors in  
 493 extreme value magnitudes. By operating in the latent space, WTL achieves superior accuracy and  
 494 sharpness in extreme weather forecasting compared to WT and other baselines like Pangu-Weather  
 495 and FourCastNet, confirming the advantage of latent space prediction for preserving extreme values.  
 496

497 **Competitive overall forecast skill.** Fig.7 (b) shows that WTL remains competitive with WT in  
 498 terms of overall forecast skill. Both models exhibit performance comparable to Pangu-Weather, and  
 499 surpass baselines like Operational IFS and FourCastNet. Notably, WTL consistently produces sharp  
 500 and accurate forecasts while adapting to multiple PVS inputs. These results suggest that conducting  
 501 weather forecasting in the latent space is a viable and efficient strategy: it matches the performance  
 502 of pixel-space models while substantially reducing data storage and computational costs.  
 503

#### 502 4.7 DOWNSTREAM TASK VERIFICATION: PRECIPITATION FORECASTING

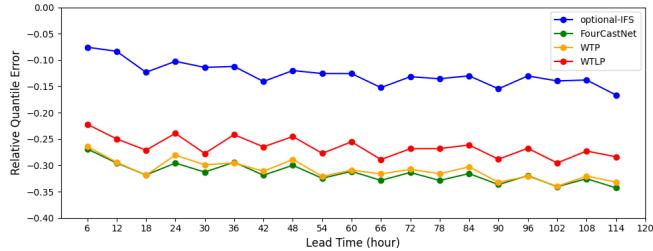
504 To further substantiate the utility  
 505 of our proposed framework beyond  
 506 general weather forecasting , we  
 507 conduct additional evaluations on  
 508 the precipitation forecasting task.  
 509 Precipitation forecasting presents a  
 510 unique challenge compared to stan-  
 511 dard atmospheric variables due to  
 512 its probability distribution, which  
 513 typically peaks strongly at zero and  
 514 exhibits a heavy tail towards pos-  
 515 itive values (Kurth et al., 2023). Ad-  
 516 dressing this challenge, we adopt  
 517 the post-processing methodology suggested by Kurth et al. (2023); Zhou et al. (2022). Specifically,  
 518 we construct a lightweight mapping network to project the 6-hour forecast states of the weather  
 519 prediction model onto 6-hour accumulated precipitation (TP6h).

520 Following the experimental protocol in Kurth et al. (2023), we employ the RQE metric and  
 521 ERA5 dataset to assess the model. We compare two variants of our framework: the Window-  
 522 ViT-Precipitation (WTP) in the pixel space, and the Window-ViT-Latent-Precipitation (WTLP) in the  
 523 latent space.

524 The experimental results in Figure 9 shows that the performance of the pixel-based WTP is comparable  
 525 to that of FourCastNet (Kurth et al., 2023), and the latent-based WTLP consistently outperforms  
 526 both WTP and FourCastNet. It demonstrates a stronger capability in characterizing the heavy-tailed  
 527 distribution of precipitation data in the latent space. While WTLP shows improvements over other  
 528 baselines, its performance remains lower than that of Operational-IFS. This highlights that accurate  
 529 extreme precipitation forecasting remains a significant challenge for purely data-driven models and  
 530 warrants further investigation.

## 531 5 CONCLUSION

532 We presented the Weather Latent Autoencoder, a novel method for learning efficient latent rep-  
 533 resentations of weather data. WLA circumvents key issues of pixel-space approaches, including  
 534 prediction smoothness and inaccuracy, single PVS limitations, and prohibitive costs. By separating  
 535 task modeling from reconstruction and unifying PVS representations, WLA facilitates accurate, PVS-  
 536 agnostic weather predictions with high efficiency. Our resulting ERA5-Latent dataset compresses  
 537 ERA5 data significantly (from 244.34 TB to 0.43 TB). WLA and the ERA5-Latent dataset offer a  
 538 robust foundation for advancing meteorological research within latent space. Future work will target  
 539 improved reconstruction and application to higher-resolution weather datasets.



533 Figure 9: Comparison on the precipitation forecasting task.  
 534 The smaller the RQE is than 0, the worse the performance in  
 535 precipitation forecasting.

540 6 STATEMENT  
541542 6.1 ETHICS STATEMENT  
543

544 The development and training phases of the presented models necessitate substantial computational  
545 resources, inherently leading to significant energy consumption. This energy expenditure constitutes  
546 a critical environmental concern, contributing materially to the carbon footprint and other associated  
547 ecological impacts. Recognizing these externalities, we emphasize the importance of mitigating  
548 strategies focused on energy sourcing. Specifically, transitioning towards renewable and low-carbon  
549 energy infrastructure for powering computational tasks is paramount to lessening the environmental  
550 burden associated with large-scale model training.

551 Adopting sustainable energy solutions can demonstrably reduce the ecological ramifications of the  
552 computational pipeline, aligning technological advancement with environmental stewardship. It  
553 is incumbent upon the research community to proactively evaluate and address the environmental  
554 costs inherent in deploying computationally demanding methodologies. Key mitigation approaches  
555 encompass not only the adoption of sustainable energy but also advancements in energy-efficient  
556 hardware architectures and the continuous pursuit of algorithmic optimization to reduce computational  
557 overhead. Promoting such holistic sustainable computational practices is crucial for ensuring that  
558 progress in artificial intelligence does not inadvertently exacerbate environmental challenges, but  
559 rather contributes responsibly to future development.

560 6.2 REPRODUCIBILITY STATEMENT  
561

562 To facilitate reproducibility, we provide comprehensive training details for the weather latent autoen-  
563 coder in the supplementary materials. Code, ERA5-latent data, and pre-trained models are available  
564 at WLA.

565 566 REFERENCES  
567

568 569 Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Omnisat: Self-supervised  
570 571 modality fusion for earth observation, 2024. URL <https://arxiv.org/abs/2404.08351>.

572 573 Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression,  
574 2017. URL <https://arxiv.org/abs/1611.01704>.

575 576 Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-  
577 578 range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.

579 580 Philippe Bougeault, Zoltan Toth, Craig Bishop, Barbara Brown, David Burridge, De Hui Chen, Beth  
581 582 Ebert, Manuel Fuentes, Thomas M Hamill, Ken Mylne, et al. The thorpx interactive grand global  
583 584 ensemble. *Bulletin of the American Meteorological Society*, 91(8):1059–1072, 2010.

585 586 Kang Chen, Tao Han, Junchao Gong, Lei Bai, Fenghua Ling, Jing-Jia Luo, Xi Chen, Leiming Ma,  
587 588 Tianning Zhang, Rui Su, Yuanzheng Ci, Bin Li, Xiaokang Yang, and Wanli Ouyang. Fengwu:  
589 590 Pushing the skillful global medium-range weather forecast beyond 10 days lead, 2023a. URL  
<https://arxiv.org/abs/2304.02948>.

591 592 Lei Chen, Xiaohui Zhong, Feng Zhang, Yuan Cheng, Yinghui Xu, Yuan Qi, and Hao Li. Fuxi: A  
593 594 cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and  
595 596 Atmospheric Science*, 6(1):190, 2023b. ISSN 2397-3722. doi: 10.1038/s41612-023-00512-1.

597 598 Liangzhi Chen, Philipp Brun, Pascal Buri, Simone Faticchi, Arthur Gessler, Michael James Mc-  
599 600 Carthy, Francesca Pellicciotti, Benjamin Stocker, and Dirk Nikolaus Karger. Global increase  
601 602 in the occurrence and impact of multiyear droughts. *Science*, 387(6731):278–284, 2025.  
603 604 doi: 10.1126/science.ado4245. URL <https://www.science.org/doi/abs/10.1126/science.ado4245>.

594 Emilien Dupont, Hrushikesh Loya, Milad Alizadeh, Adam Goliński, Yee Whye Teh, and Arnaud  
 595 Doucet. Coin++: Neural compression across modalities, 2022. URL <https://arxiv.org/abs/2201.12904>.  
 596

597 Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image  
 598 synthesis, 2021. URL <https://arxiv.org/abs/2012.09841>.  
 599

600 Junchao Gong, Lei Bai, Peng Ye, Wanghan Xu, Na Liu, Jianhua Dai, Xiaokang Yang, and Wanli  
 601 Ouyang. Cascast: Skillful high-resolution precipitation nowcasting via cascaded modelling, 2024a.  
 602 URL <https://arxiv.org/abs/2402.04290>.  
 603

604 Junchao Gong, Siwei Tu, Weidong Yang, Ben Fei, Kun Chen, Wenlong Zhang, Xiaokang Yang, Wanli  
 605 Ouyang, and Lei Bai. Postcast: Generalizable postprocessing for precipitation nowcasting via  
 606 unsupervised blurriness modeling, 2024b. URL <https://arxiv.org/abs/2410.05805>.  
 607

608 Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu.  
 609 Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis, 2025. URL  
 610 <https://arxiv.org/abs/2412.04431>.  
 611

612 Tao Han, Zhenghao Chen, Song Guo, Wanghan Xu, and Lei Bai. Cra5: Extreme compression of era5  
 613 for portable global climate and weather research via an efficient variational transformer, 2024a.  
 614 URL <https://arxiv.org/abs/2405.03376>.  
 615

616 Tao Han, Song Guo, Fenghua Ling, Kang Chen, Junchao Gong, Jingjia Luo, Junxia Gu, Kan Dai,  
 617 Wanli Ouyang, and Lei Bai. Fengwu-ghr: Learning the kilometer-scale medium-range global  
 618 weather forecasting, 2024b.  
 619

620 Philippe Hansen-Estruch, David Yan, Ching-Yao Chung, Orr Zohar, Jialiang Wang, Tingbo Hou, Tao  
 621 Xu, Sriram Vishwanath, Peter Vajda, and Xinlei Chen. Learnings from scaling visual tokenizers  
 622 for reconstruction and generation, 2025. URL <https://arxiv.org/abs/2501.09755>.  
 623

624 Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. Elic: Efficient learned  
 625 image compression with unevenly grouped space-channel contextual adaptive coding. In 2022  
 626 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5708–5717,  
 627 2022. doi: 10.1109/CVPR52688.2022.00563.  
 628

629 Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hihara, András Horányi, Joaquín Muñoz-Sabater,  
 630 Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci,  
 631 Saleh Abdalla, Xavier Abellán, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot,  
 632 Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana  
 633 Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger,  
 634 Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux,  
 635 Philippe Lopez, Cristina Lupu, Gabor Radnoti, Patricia De Rosnay, Iryna Rozum, Freja Vamborg,  
 636 Sébastien Villaume, and Jean-Noël Thépaut. The era5 global reanalysis. *Quarterly Journal of the  
 637 Royal Meteorological Society*, 146(730):1999–2049, 2020. ISSN 0035-9009, 1477-870X. doi:  
 638 10.1002/qj.3803.  
 639

640 Tian Hua and Li Chong-Yin. Further study of typhoon tracks and the low-frequency (30-60 days)  
 641 wind-field pattern at 850 hpa. *Atmospheric and Oceanic Science Letters*, 3(6):319–324, 2010.  
 642 ISSN 1674-2834, 2376-6123. doi: 10.1080/16742834.2010.11446891.  
 643

644 Langwen Huang and Torsten Hoefer. Compressing multidimensional weather and climate data into  
 645 neural networks, 2023. URL <https://arxiv.org/abs/2210.12538>.  
 646

647 A. Inness, M. Ades, A. Agustí-Panareda, J. Barré, A. Benedictow, A.-M. Blechschmidt, J. J.  
 648 Dominguez, R. Engelen, H. Eskes, J. Flemming, V. Huijnen, L. Jones, Z. Kipling, S. Massart,  
 649 M. Parrington, V.-H. Peuch, M. Razinger, S. Remy, M. Schulz, and M. Suttie. The cams reanalysis  
 650 of atmospheric composition. *Atmospheric Chemistry and Physics*, 19(6):3515–3556, 2019.  
 651 doi: 10.5194/acp-19-3515-2019. URL <https://acp.copernicus.org/articles/19/3515/2019/>.  
 652

653 Milan Klöwer, Miha Razinger, Juan J. Dominguez, Peter D. Düben, and Tim N. Palmer. Compressing  
 654 atmospheric data into its real information content. *Nature Computational Science*, 1(11):713–724,  
 655 2021. ISSN 2662-8457. doi: 10.1038/s43588-021-00156-2.  
 656

648 Robert J. Kuligowski and Ana P. Barros. Experiments in short-term precipitation  
 649 forecasting using artificial neural networks. *Monthly Weather Review*, 126(2):470  
 650 – 482, 1998. doi: 10.1175/1520-0493(1998)126<0470:EISTPF>2.0.CO;2. URL  
 651 [https://journals.ametsoc.org/view/journals/mwre/126/2/1520-0493\\_1998\\_126\\_0470\\_eistpf\\_2.0.co\\_2.xml](https://journals.ametsoc.org/view/journals/mwre/126/2/1520-0493_1998_126_0470_eistpf_2.0.co_2.xml).  
 652

653 Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David  
 654 Hall, Andrea Miele, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: Accelerating global  
 655 high-resolution weather forecasting using adaptive fourier neural operators. In *Proceedings of the*  
 656 *Platform for Advanced Scientific Computing Conference*, PASC '23, New York, NY, USA, 2023.  
 657 Association for Computing Machinery. ISBN 9798400701900. doi: 10.1145/3592979.3593412.  
 658 URL <https://doi.org/10.1145/3592979.3593412>.  
 659

660 Xin Liang, Kai Zhao, Sheng Di, Sihuan Li, Robert Underwood, Ali M. Gok, Jiannan Tian, Junjing  
 661 Deng, Jon C. Calhoun, Dingwen Tao, Zizhong Chen, and Franck Cappello. Sz3: A modular  
 662 framework for composing prediction-based error-bounded lossy compressors. *IEEE Transactions*  
 663 *on Big Data*, 9(2):485–498, 2023. doi: 10.1109/TBDA.2022.3201176.

664 Nian LIU, Zhongwei YAN, Xuan TONG, Jiang JIANG, Haochen LI, Jiangjiang XIA, Xiao LOU,  
 665 Rui REN, and Yi FANG. Meshless surface wind speed field reconstruction based on machine  
 666 learning. *ADVANCES IN ATMOSPHERIC SCIENCES*, 39(10):1721–1733, 2022. ISSN 0256-1530.  
 667 doi: 10.1007/s00376-022-1343-8. URL <http://www.iapjournals.ac.cn/aas/en/article/doi/10.1007/s00376-022-1343-8>.  
 668

669 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL <https://arxiv.org/abs/1711.05101>.  
 670

672 Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical  
 673 full resolution learned lossless image compression. In *2019 IEEE/CVF Conference on Computer*  
 674 *Vision and Pattern Recognition (CVPR)*, pp. 10621–10630, 2019. doi: 10.1109/CVPR.2019.01088.  
 675

676 Piotr Mirowski, David Warde-Farley, Mihaela Rosca, Matthew Koichi Grimes, Yana Hasson, Hyunjik  
 677 Kim, Mélanie Rey, Simon Osindero, Suman Ravuri, and Shakir Mohamed. Neural compression of  
 678 atmospheric states, 2024. URL <https://arxiv.org/abs/2407.11666>.  
 679

680 Todd W. Moore and Richard W. Dixon. Patterns in 500 hpa geopotential height associated with  
 681 temporal clusters of tropical cyclone tornadoes: 500 hpa geopotential height patterns and tropical  
 682 cyclone tornadoes. *Meteorological Applications*, 22(3):314–322, 2015. ISSN 13504827. doi:  
 683 10.1002/met.1456.

684 Jonathan A. Patz, Diarmid Campbell-Lendrum, Tracey Holloway, and Jonathan A. Foley. Impact of  
 685 regional climate change on human health. *Nature*, 438(7066):310–317, 2005. ISSN 0028-0836,  
 686 1476-4687. doi: 10.1038/nature04188.

687 Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R. Andersson, Andrew El-Kadi, Dominic  
 688 Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, Remi Lam, and Matthew  
 689 Willson. Gencast: Diffusion-based ensemble forecasting for medium-range weather, 2024. URL  
 690 <https://arxiv.org/abs/2312.15796>.  
 691

692 Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukaina Mouatadid, and  
 693 Nils Thuerey. Weatherbench: A benchmark data set for data-driven weather forecasting. *Journal*  
 694 *of Advances in Modeling Earth Systems*, 12(11):e2020MS002203, 2020. doi: <https://doi.org/10.1029/2020MS002203>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020MS002203>. e2020MS002203 10.1029/2020MS002203.  
 695

696 Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan  
 697 Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, Rachel Prudden, Amol Mand-  
 698 hane, Aidan Clark, Andrew Brock, Karen Simonyan, Raia Hadsell, Niall Robinson, Ellen Clancy,  
 699 Alberto Arribas, and Shakir Mohamed. Skilful precipitation nowcasting using deep genera-  
 700 tive models of radar. *Nature*, 597(7878):672–677, 2021. ISSN 0028-0836, 1476-4687. doi:  
 701 10.1038/s41586-021-03854-z.

702 Fuyou Tian, Yongguang Zheng, Tao Zhang, Xiaoling Zhang, Dongyan Mao, Jianhua Sun, and Sixiong  
 703 Zhao. Statistical characteristics of environmental parameters for warm season short-duration heavy  
 704 rainfall over central and eastern china. *Journal of Meteorological Research*, 29(3):370–384, 2015.  
 705 ISSN 2095-6037, 2198-0934. doi: 10.1007/s13351-014-4119-y.

706 Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning.  
 707 In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and  
 708 R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Asso-  
 709 ciates, Inc., 2017. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/7a98af17e63a0ac09ce2e96d03992fbc-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/7a98af17e63a0ac09ce2e96d03992fbc-Paper.pdf).

710 Kristoffer H. Wild, Raymond B. Huey, Eric R. Pianka, Susana Clusella-Trullas, Anthony L. Gilbert,  
 711 Donald B. Miles, and Michael R. Kearney. Climate change and the cost-of-living squeeze  
 712 in desert lizards. *Science*, 387(6731):303–309, 2025. doi: 10.1126/science.adq4372. URL  
 713 <https://www.science.org/doi/abs/10.1126/science.adq4372>.

714 Yueqi Xie, Ka Leong Cheng, and Qifeng Chen. Enhanced invertible encoding for learned im-  
 715 age compression. In *Proceedings of the 29th ACM International Conference on Multimedia*,  
 716 MM ’21, pp. 162–170, New York, NY, USA, 2021. Association for Computing Machinery.  
 717 ISBN 9781450386517. doi: 10.1145/3474085.3475213. URL <https://doi.org/10.1145/3474085.3475213>.

718 Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J. Stewart, Joëlle Hanna, Damian Borth, Ioannis  
 719 Papoutsis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired  
 720 multimodal foundation model for earth observation, 2024. URL <https://arxiv.org/abs/2403.15356>.

721 Wanghan Xu, Kang Chen, Tao Han, Hao Chen, Wanli Ouyang, and Lei Bai. Extremecast: Boosting  
 722 extreme value prediction for global weather forecast, 2024. URL <https://arxiv.org/abs/2402.01295>.

723 Dangfu Yang, Shengjun Liu, Yamin Hu, Xinru Liu, Jiehong Xie, and Liang Zhao. Predictor selection  
 724 for cnn-based statistical downscaling of monthly precipitation. *Advances in Atmospheric Sciences*,  
 725 40(6):1117–1131, 2023. ISSN 0256-1530, 1861-9533. doi: 10.1007/s00376-022-2119-x.

726 Keer Zhang, Chang Cao, Haoran Chu, Lei Zhao, Jiayu Zhao, and Xuhui Lee. Increased heat risk in  
 727 wet climate induced by urban humid heat. *Nature*, 617(7962):738–742, 2023a.

728 Yuchen Zhang, Mingsheng Long, Kaiyuan Chen, Lanxiang Xing, Ronghua Jin, Michael I. Jordan,  
 729 and Jianmin Wang. Skilful nowcasting of extreme precipitation with nowcastnet. *Nature*, 619  
 730 (7970):526–532, 2023b. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-023-06184-4.

731 Yue Zhao, Yuanjun Xiong, and Philipp Krähenbühl. Image and video tokenization with binary  
 732 spherical quantization, 2024. URL <https://arxiv.org/abs/2406.07548>.

733 Kanghui Zhou, Jisong Sun, Yongguang Zheng, and Yutao Zhang. Quantitative precipitation forecast  
 734 experiment based on basic nwp variables using deep learning. *Advances in Atmospheric Sciences*,  
 735 39(9):1472–1486, 2022.

736  
 737  
 738  
 739  
 740  
 741  
 742  
 743  
 744  
 745  
 746  
 747  
 748  
 749  
 750  
 751  
 752  
 753  
 754  
 755

756 A SUPPLEMENTARY MATERIAL  
757758 A.1 EXPERIMENTAL DETAILS OF WEATHER LATENT AUTOENCODER  
759760 A.1.1 ERA5 DATASET DETAILS  
761

762 The ERA5 dataset (Hersbach et al., 2020) is a global atmospheric reanalysis product from the  
763 European Center for Medium-Range Weather Forecasts (ECMWF) and serves as a standard for  
764 evaluating the Weather Latent Autoencoder in comparison to other models. This dataset is highly  
765 valued in climate research due to its high spatial resolution of  $0.25^\circ$  and extensive weather coverage.  
766 The dataset is temporally partitioned into training sets (1979–2021, 233.48 TB), validation sets (2022,  
767 5.43 TB), and test sets (2023, 5.43 TB). To meet various weather-related needs, we organized three  
768 categories of variables within the ERA5 dataset: upper-air, surface, and precipitation variables.  
769

770 For **upper-air variables**, we selected three configurations spanning 25, 13, and 6 pressure levels,  
771 each containing six core weather variables: geopotential height ( $z$ ), longitudinal wind speed ( $u$ ),  
772 meridional wind speed ( $v$ ), vertical velocity ( $w$ ), atmospheric temperature ( $t$ ), and specific humidity  
773 ( $q$ ), in which variables are presented by abbreviating their short name and pressure levels (e.g.,  $q1000$   
774 denotes the specific humidity at a pressure level of 1000 hPa).

775 The **surface variables** include two subsets: an 8-variable set comprising 10m v-component of wind  
776 (10v), 10m u-component of wind (10u), 100m v-component of wind (100v), 100m u-component of  
777 wind (100u), 2m temperature (t2m), Total cloud cover (tcc), surface pressure (sp) and Mean sea-level  
778 pressure (msl); and a streamlined 4-variable subset (10v, 10u, tcc, msl).

779 The **precipitation variables** cover cumulative hourly precipitation over six intervals (tp1h, tp2h,  
780 tp3h, tp4h, tp5h, tp6h), with two additional single-variable subsets (tp1h and tp6h).

781 A.1.2 IMPLEMENTATION DETAILS  
782

783 To accommodate the distinct physical properties of atmospheric data, we employ a variable-specific  
784 training strategy. Specifically, we train separate WLAs for each of the six upper-air variables, and  
785 distinct models for the surface and precipitation variables. This design treats different variables as  
786 independent modalities (Chen et al., 2023a), ensuring that the unique physical laws governing each  
787 variable do not interfere during feature extraction, thereby maximizing reconstruction fidelity.  
788

789 All WLAs are trained with identical configurations across weather variables. The models are  
790 optimized for 500K steps on 4 Tesla A100 GPUs. The codebook size is set to  $2^{128}$  for upper-air and  
791 surface variables, while precipitation variables employ a reduced codebook size of  $2^{32}$  due to their  
792 higher compressibility. The input data are processed through patches of size  $15 \times 14$  with a stride of  
793  $10 \times 10$  and a padding of  $2 \times 2$ .

794 Following the architectural insights of (Hansen-Estruch et al., 2025), where decoder upscaling  
795 demonstrated significant reconstruction benefits without comparable encoder gains, we design the  
796 VAEformer with asymmetric depths: a 16-layer encoder versus a 32-layer decoder. Training employs  
797 the AdamW (Loshchilov & Hutter, 2019) optimizer with an initial learning rate of  $3.2 \times 10^{-5}$ , batch  
798 size 8, and a hybrid learning schedule that combines a linear warm-up phase increasing the learning  
799 rate from  $3.2 \times 10^{-6}$  to  $3.2 \times 10^{-5}$ , followed by a cosine decay phase.  
800

801 Following the settings in BSQ (Zhao et al., 2024), the total loss function comprises three parts:  
802 entropy loss, MSE loss, and GAN loss. Among them, entropy loss is used to improve the utilization  
803 rate of the codebook, and MSE loss and GAN loss are used to improve the accuracy and clarity of the  
804 reconstruction results.  
805

806 A.1.3 BASELINE  
807

808 The compared state-of-the-art compression methods include Elic (He et al., 2022), IEN (Xie et al.,  
809 2021), VQVAE (van den Oord et al., 2017), VQGAN (Esser et al., 2021), and VAEformer (Han et al.,  
810 2024a). Elic (He et al., 2022) introduces an efficient architecture that utilizes unevenly grouped  
811 space-channel contextual adaptive coding, striking a balance between rate-distortion performance and  
812 computational complexity. IEN (Xie et al., 2021) enhances the compression pipeline by proposing  
813 a more powerful invertible encoding network, which improves the modeling of the latent repre-  
814

810 representations' distribution. The comparison also includes methods based on discrete representations.  
 811 VQ-VAE (van den Oord et al., 2017) is a foundational model that learns a discrete codebook for latent  
 812 variables, effectively preventing posterior collapse. Building upon this, VQGAN (Esser et al., 2021)  
 813 combines the discrete quantization of VQ-VAE with the high-fidelity synthesis power of Transformers  
 814 and GANs. Lastly, VAEformer (Han et al., 2024a) adapts a variational transformer architecture  
 815 specifically for the extreme compression of large-scale scientific data, leveraging the transformer's  
 816 ability to capture complex long-range dependencies.

817 For the Elic (He et al., 2022), IEN (Xie et al., 2021), we use the code in CompressAI 2 (Ballé et al.,  
 818 2017) to reimplement and retrain them. For the VQVAE (van den Oord et al., 2017), VQGAN (Esser  
 819 et al., 2021), we fine-tune their pre-trained models on meteorological data. For VAEformer (Han  
 820 et al., 2024a), since its pre-trained model uses data consistent with this study, we directly use its  
 821 pre-trained model for comparison.

822 The most challenging baseline among all baselines is VAEFormer. VAEFormer is specifically  
 823 designed for meteorological data compression. It uses the atmospheric circulation transformer block  
 824 as a basic block to effectively capture the characteristics of atmospheric circulation. Meanwhile,  
 825 VAEFormer includes two stages: pre-training and fine-tuning. The pre-training stage trains a VAE-  
 826 style transformer encoder that generates the compressed latent representation, and a transformer-based  
 827 decoder restores it to the reconstructed data. The fine-tuning stage trains another encoder and decoder  
 828 to predict the mean and scale hyperpriors for the Arithmetic Encoder and Decoder process, which  
 829 further losslessly compresses the data via entropy coding.

## 830 A.2 EXPERIMENTAL DETAILS OF DOWNSTREAM

### 831 A.2.1 ERA5-LATENT DATASET DETAILS

832 The partitioning of ERA5-Latent data and the selection of PVS remain consistent with Section 4.1.  
 833 By utilizing the high compression rate of WLA, the ERA5-Latent dataset reduces the original 244.34  
 834 TB of data down to 0.43 TB, while providing a unified representation for multiple PVS. These subsets  
 835 include three for upper-air variables corresponding to 25, 13, and 6 pressure levels, two for surface  
 836 variables (4 and 8 variables), and three for precipitation variables (tp1h, tp6h, tp1-6h). To facilitate  
 837 the computation of pixel metrics, the ERA5-Latent dataset also incorporates the raw pixel data for  
 838 July 2023, which has been compressed using the Lempel-Ziv-Markov chain-Algorithm (LZMA) and  
 839 occupies 0.117 TB of storage. LZMA is a widely used lossless compression algorithm developed by  
 840 Igor Pavlov and implemented via the Python standard library.

841 Building upon the ERA5-Latent dataset, deep learning models for large-scale meteorological research  
 842 can utilize the unified latent representation to seamlessly handle multiple PVS, making the models  
 843 adaptable for a wide range of meteorological tasks and scenarios. Moreover, processes that typically  
 844 require large datasets, such as training, validation, and testing, can be conducted using the compact  
 845 latent data, significantly reducing both storage and data computational costs. For tasks requiring only  
 846 a limited amount of pixel data, pixel data from a single month within the ERA5-Latent dataset can be  
 847 used to compute metrics and visually compare model outputs with the original data.

### 851 A.2.2 MODEL DETAILS

852 Window-ViT employs the multimodal encoder-decoder from FengWu (Chen et al., 2023a) and the  
 853 backbone of the weather latent autoencoder, conducting weather forecasting task in the pixel space.  
 854 It treats each upper-air variable and the grouped surface variables as distinct modalities, and uses  
 855 local and global self-attention to model complex atmospheric dynamics. In contrast, Window-ViT-  
 856 Latent omits the downsampling and upsampling components present in WT while keeping all other  
 857 components unchanged, conducting weather forecasting task in the latent space.

### 860 A.2.3 VARIABLE AND METRIC DETAILS

861 The variables used in the weather forecasting experiment are consistent with the settings of PanGu-  
 862 Weather, including five upper-air variables (z, u, v, t, q) and four surface variables (10v, 10u, t2m,  
 863 msl).

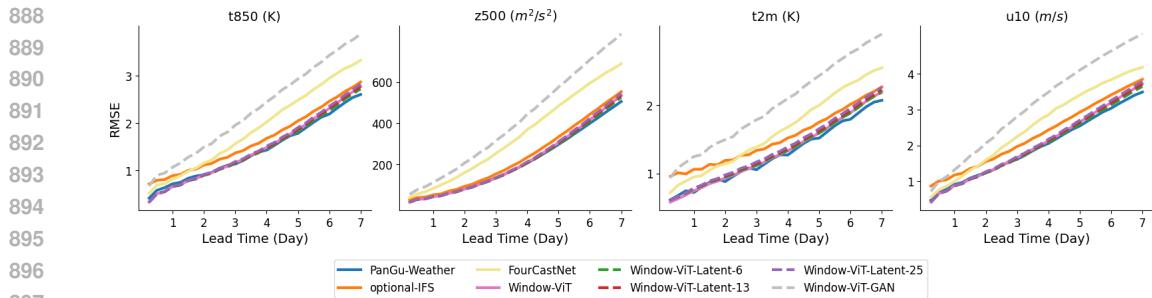
864 The SEDI metric classifies each pixel into extreme or normal weather using high quantile thresholds  
 865 (90%, 95%, 98%, and 99%) and then calculates the hit rate, a value closer to 1 indicates a more  
 866 accurate prediction of extreme weather.

867 The RQE (Kurth et al., 2023) is an indicator used to assess a model’s ability to capture extreme  
 868 values within a given field, such as wind speed or precipitation. Its calculation involves summing up  
 869 the relative difference between the predicted and true values across a range of high quantiles (like  
 870 the 90th to 99.99th percentiles), focusing on the most extreme events. Essentially, the RQE reveals  
 871 systematic biases: a negative RQE suggests the model is consistently under-predicting the magnitude  
 872 of these extremes, while a positive RQE would indicate over-prediction.

#### 874 A.2.4 COMPARED MODELS

875 The optional Integrated Forecasting System of the European Centre for Medium-Range Weather  
 876 Forecasts (ECMWF) is widely regarded as the world’s leading global numerical weather prediction  
 877 (NWP) system (Bougeault et al., 2010). It utilizes a comprehensive Earth system model and an  
 878 advanced data assimilation system to produce forecasts for the medium-range and beyond. The IFS  
 879 simulates the complex interactions within the Earth’s atmosphere and its coupled systems, solving  
 880 mathematical equations that govern their dynamics and physics to predict future weather conditions.

881 PanGu-Weather is a deep learning-based weather forecasting system (Bi et al., 2023). It is a data-  
 882 driven model that employs a 3D deep neural network architecture to capture intricate patterns in  
 883 atmospheric data. Trained on decades of global reanalysis data, PanGu-Weather demonstrates strong  
 884 performance in medium-range forecasting for various atmospheric variables. Unlike traditional NWP  
 885 models, it does not explicitly solve physical equations but learns the evolution of weather patterns  
 886 directly from historical data.



888  
 889 Figure S1: Comparison on the weather forecasting task over a 7-day period. The number following  
 890 “Window-ViT-latent” indicates the number of input pressure levels.

#### 902 A.3 IMPACT OF ADVERSARIAL LOSS

903 To address the concern regarding whether the improved sharpness in extreme event prediction is  
 904 merely a methodological artifact of the loss function rather than a benefit of our proposed latent  
 905 framework, we conducted an ablation study comparing our approach against a pixel-space baseline  
 906 trained with adversarial loss.

907 We trained the baseline pixel-space WT with an additional Generative Adversarial Network (GAN)  
 908 loss, referred to as Window-ViT-GAN (WTG). All other hyperparameters and training configurations  
 909 remained identical to the standard WT baseline to ensure a fair comparison.

910 The results in Figure S1 show that the WTG model failed to surpass the standard baselines. Instead of  
 911 improving extreme event prediction, the instability compromised the model’s ability to learn correct  
 912 temporal dynamics, resulting in inferior predictive performance across standard metrics compared  
 913 to both the vanilla WT and other baselines. These results underscore the non-triviality of applying  
 914 adversarial training to global weather forecasting. Merely altering the loss function in pixel space  
 915 leads to optimization difficulties. This validates the design of our Latent Framework, which effectively  
 916 decouples the objectives: the WLA utilizes GAN loss to ensure sharp reconstruction, while the WTL  
 917 employs Binary Cross-Entropy loss to stably learn temporal dynamics in the latent space.

918 A.4 LIMITATIONS AND EXPECTATIONS  
919920 Although the weather latent autoencoder effectively transforms weather data from pixel space to  
921 latent space, we acknowledge certain limitations, particularly regarding the computational trade-offs  
922 and scalability.923 First, to accommodate the distinct physical properties of atmospheric variables and ensure high-fidelity  
924 reconstruction, we trained separate autoencoders for distinct variable groups (upper-air, surface, and  
925 precipitation). This approach is both data-intensive and computationally expensive. Consequently,  
926 both the model training and the generation of the comprehensive ERA5-Latent dataset entail a  
927 substantial upfront computational investment. We frame this as a necessary trade-off: a significant  
928 one-time cost is incurred to enable substantial long-term reductions in storage and computational  
929 overhead for downstream applications and the broader research community.930 Second, the current architecture is primarily optimized for ERA5 data at a resolution of  $0.25^\circ$ . Scaling  
931 to higher-resolution datasets (e.g., HRES at  $0.09^\circ$ ) remains a challenge, as the exceedingly large  
932 global weather images impose significant computational demands during training.933 In future work, we plan to build upon the foundation of the ERA5-based autoencoder and employ  
934 efficient fine-tuning techniques to adapt the model to both global and regional datasets at higher  
935 spatial resolutions. This strategy aims to further mitigate computational barriers and facilitate a  
936 paradigm shift towards conducting scalable weather research within the latent space.937  
938 A.5 VISUALIZATION OF RECONSTRUCTED RESULTS  
939940 To provide a detailed assessment of our model’s reconstruction fidelity and to transparently analyze  
941 potential information loss due to compression, this section presents a comprehensive visual analysis.  
942 We examine six key variables from the ERA5 dataset: Atmospheric temperature at 850hpa (T850),  
943 Geopotential height at 500hpa (Z500), Temperature at 2 meters (t2m), U-component of wind at 10  
944 meters (10u), Mean sea-level pressure (msl), and 6-hour accumulated precipitation (tp6h).945 As illustrated in Figures S1 through S6, we present three case studies for each variable. Each case  
946 includes a side-by-side comparison of: (a) the original field from ERA5, (b) the field reconstructed  
947 from our model’s latent representation, and (c) a difference map (i.e., reconstruction error) to precisely  
948 identify the magnitude and location of any information loss.949 The reconstruction and difference maps demonstrate that our model achieves superior reconstruction  
950 quality across this diverse set of atmospheric variables. Critically, the model proves highly effective at  
951 preserving the fine-grained, small-scale extreme values that are vital for meteorological applications.  
952 For instance, the visualizations confirm the retention of sharp gradients in Z500 fields associated  
953 with atmospheric troughs and ridges, the intricate structures of precipitation bands (tp6h) within  
954 storm systems, and the tight pressure contours (msl) that define the core of cyclones. This robust  
955 performance underscores the model’s ability to effectively capture and represent the essential features  
956 of the original data.957 While the overall reconstruction fidelity is high, a closer examination of the difference maps reveals  
958 that some localized information loss does occur. This loss is not random but is systematically  
959 correlated with regions of extreme values in the original data. The spatial patterns of this information  
960 loss are variable-dependent, as the distribution of extreme values is intrinsically tied to the physical  
961 nature of each atmospheric field. For example:962  
963 1. For 6-hour accumulated precipitation (tp6h), information loss predominantly manifests in  
964 localized cores of intense convective rainfall, such as within tropical cyclones or severe  
965 thunderstorms.  
966  
967 2. For temperature at 2 meters (t2m), larger reconstruction errors are more likely to appear in  
968 areas with extreme temperatures, such as the polar regions or hot deserts.  
969  
970 3. For geopotential height at 500hpa (Z500), the most notable discrepancies are found at the  
971 centers of deep, low-pressure troughs or high-pressure ridges, which represent the maxima  
and minima of the atmospheric wave patterns.

This detailed visual analysis provides a transparent and nuanced understanding of our model's performance, confirming its high fidelity while also characterizing the predictable and physically grounded nature of the minor information loss.

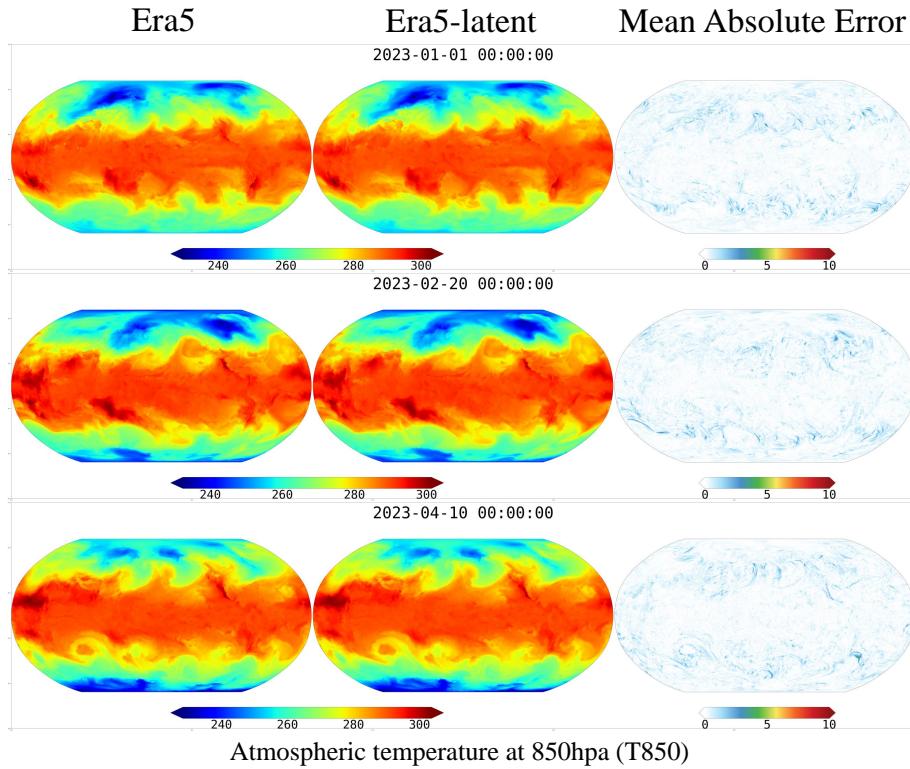


Figure S2: Visualization samples of T850 on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039

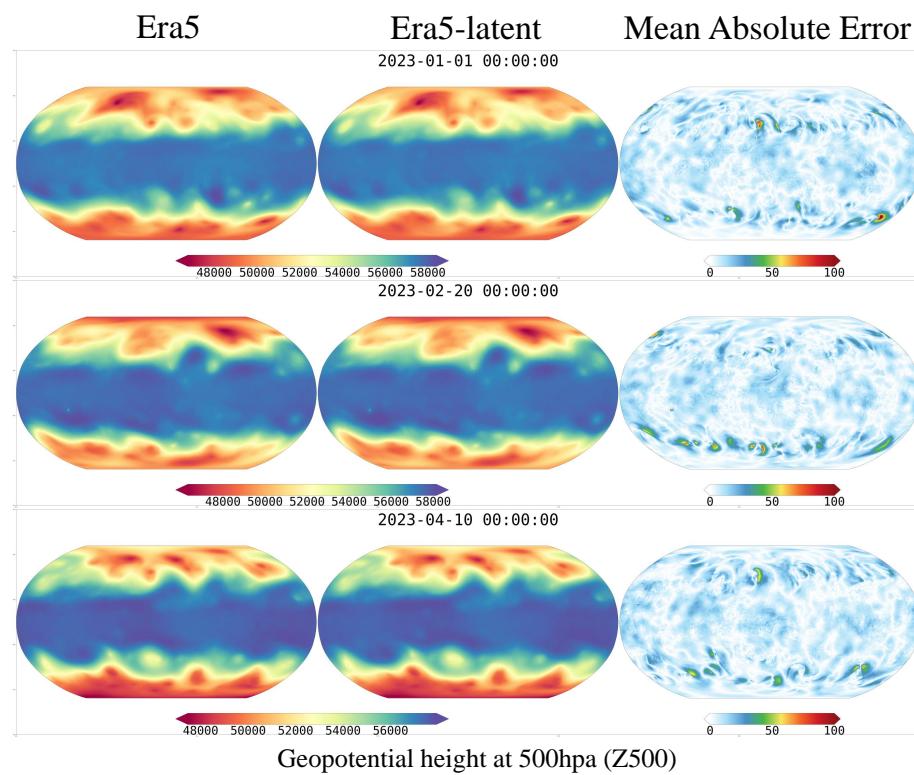


Figure S3: Visualization samples of Z500 on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

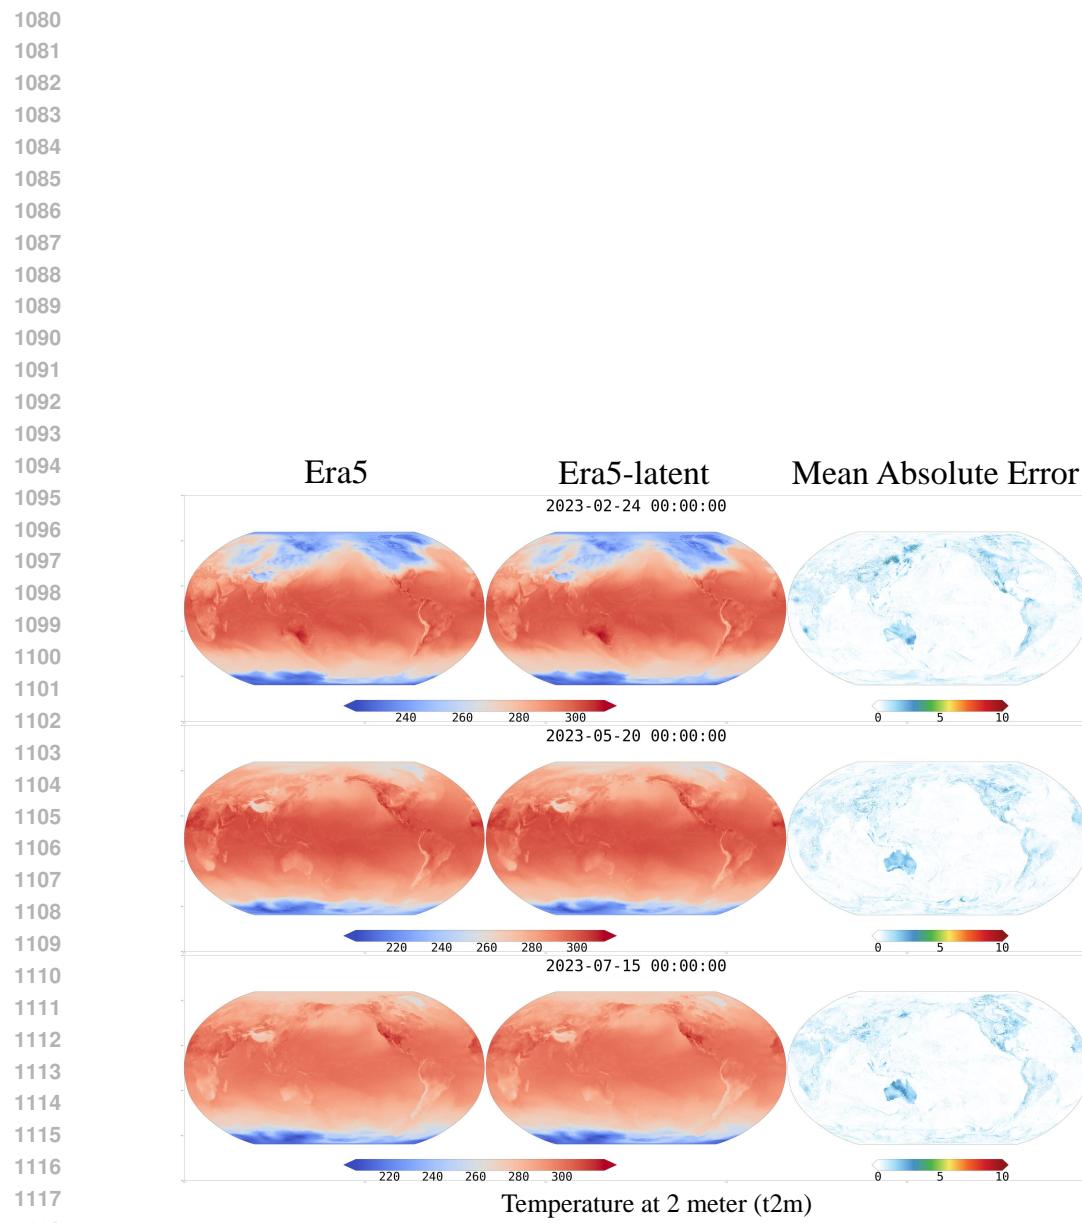


Figure S4: Visualization samples of t2m on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147

1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172

1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187

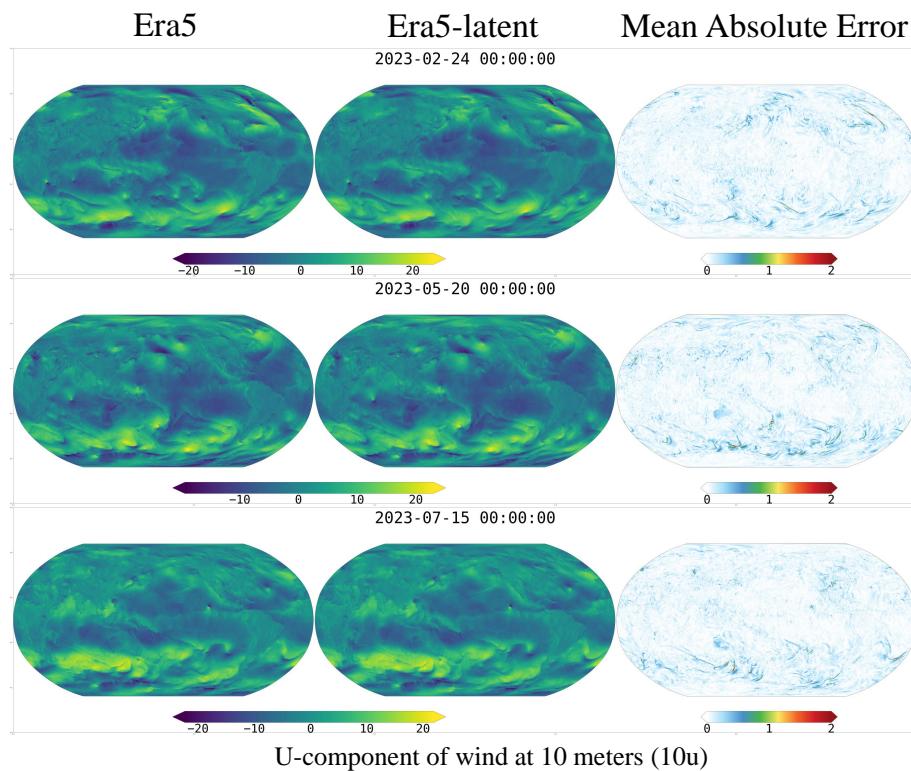


Figure S5: Visualization samples of 10u on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201

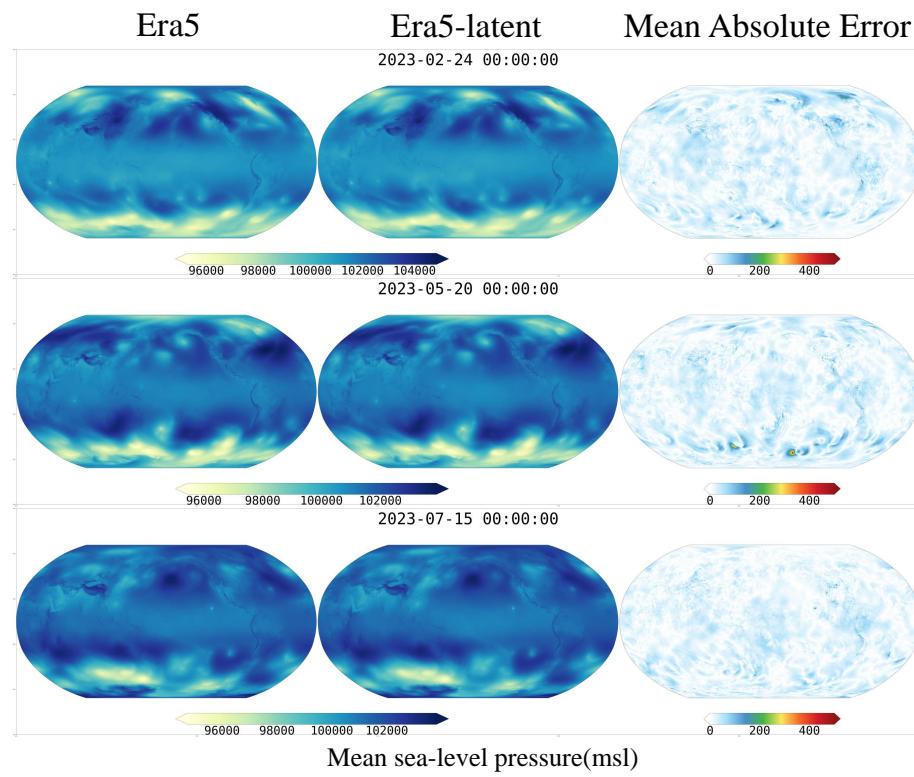


Figure S6: Visualization samples of msl on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

1242  
1243  
1244  
1245  
1246  
1247  
1248  
1249  
1250  
1251  
1252  
1253  
1254  
1255

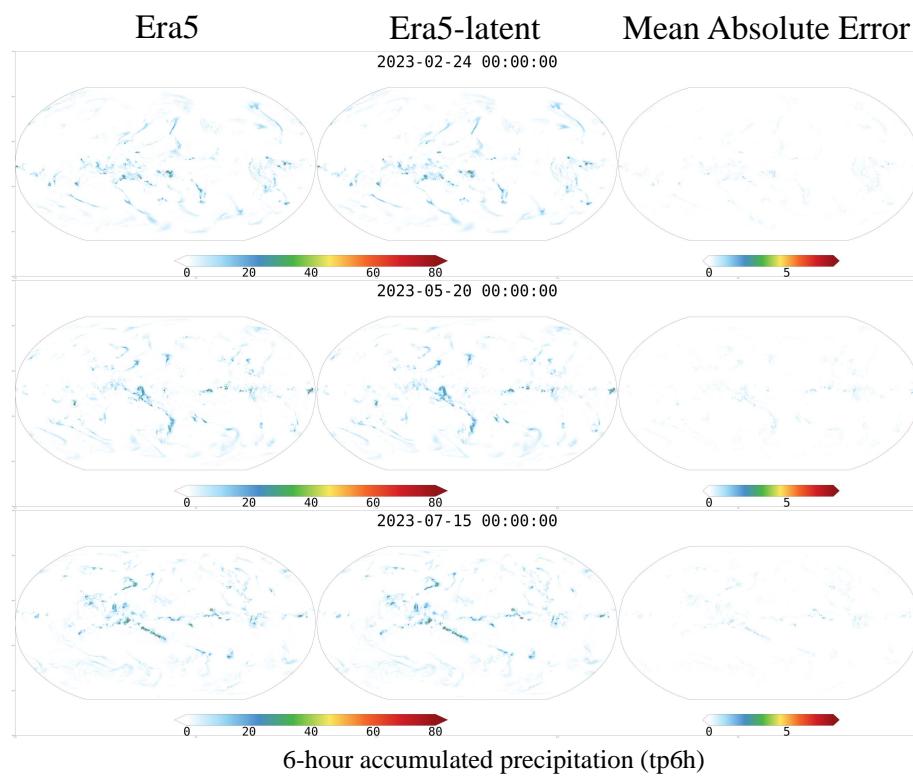


Figure S7: Visualization samples of tp6h on the ERA5 and the compressed ERA5-Latent. From the left to the right column: ERA5, ERA5-Latent, and their absolute error map.

1283  
1284  
1285  
1286  
1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295