RESEARCH ARTICLE
ATMOSPHERIC PHYSICS

Earth and Planetary Physics
10: 454–462, 2026
doi: 10.26464/epp2026046

A transfer learning framework for the real-time detection of atmospheric gravity waves from All-Sky Airglow Imager

YuBin He[1,2], QingChen Xu[1,2*], YaJun Zhu[1,2], QinZeng Li[1,2], Cui Tu[1,2], Bing Cai[1,2], Wei Yuan[1], XinYing Wang[1], and Feng Wei[1]

[1]State Key Laboratory of Solar Activity and Space Weather, National Space Science Center, Chinese Academy of Sciences, Beijing 100190, China;

[2]University of Chinese Academy of Sciences, Beijing 100049, China

Key Points:

- A novel real-time detection framework is proposed, integrating domain-knowledge-based image enhancement with the deformable convolution enhanced EfficientNet-B3 for highly robust feature extraction from All-Sky Airglow Imager.
- The framework achieves high identification accuracy of 91.2% with millisecond-level inference speed, enabling practical, real-time analysis of continuous data streams.
- This deployable solution successfully automates the climatological statistics of gravity waves, overcoming manual bottlenecks and enhancing the timeliness of space weather monitoring.

Citation: He, Y. B., Xu, Q. C., Zhu, Y. J., Li, Q. Z., Tu, C., Cai, B., Yuan, W., Wang, X. Y., and Wei, F. (2026). A transfer learning framework for the real-time detection of atmospheric gravity waves from All-Sky Airglow Imager. *Earth Planet. Phys.*, *10*(3), 454–462. http://doi.org/10.26464/epp2026046

Abstract: Atmospheric gravity waves (AGWs) observed by the All-Sky Airglow Imager (ASAI) require accurate identification for the study of atmospheric coupling mechanisms and space weather prediction. However, the traditional manual screening methods and existing machine learning approaches do not meet the demands of practical station monitoring, which has significantly impeded climatological statistical research based on AGWs. Therefore, a real-time detection framework for ground-based airglow gravity waves that integrates transfer learning with adaptive image preprocessing has been proposed. By employing wavelength-adaptive median filtering and multiscale fusion, the framework effectively suppresses stellar noise while preserving weak gravity wave features. The model utilizes an EfficientNet-B3 (convolutional neural network) backbone enhanced with a deformable convolutional layer, trained via a two-stage strategy: A frozen phase prevents overfitting by locking the lower level feature extractor, and a fine-tuning phase optimizes the deformable convolution through cosine annealing and layered optimization. This approach improves both feature transfer efficiency and gravity wave detection sensitivity. The resulting lightweight model achieves 91.2% accuracy with millisecond-level inference speed (23 ms per frame).

Keywords: atmospheric gravity waves; All-Sky Airglow Imager; real-time detection; multiscale fusion; transfer learning; EfficientNet-B3; deformable convolution

---

1. Introduction

Atmospheric gravity waves (AGWs) are a crucial wave phenomenon in Earth's atmosphere, with typical scales ranging from tens to hundreds of kilometers (Fritts, 1984; Nappo, 2002). As a primary carrier for the vertical transport of atmospheric energy and momentum, the propagation, breaking, and dissipation processes of gravity waves profoundly influence the circulation structure, temperature distribution, and chemical composition of the middle and upper atmosphere (Fritts and Alexander, 2003; Duan BQ and Pan WL, 2017). The momentum deposition of gravity waves in the mesosphere and lower thermosphere region directly affects atmospheric circulation stability and the evolution of space weather (Tu C and Hu X, 2012). Therefore, an understanding of the excitation mechanisms, propagation characteristics, and energy transport processes of gravity waves holds significant importance for revealing atmospheric coupling mechanisms, improving global climate models, and enhancing the accuracy of space weather forecasting (Ayorinde et al., 2024).

The ground-based All-Sky Airglow Imager (ASAI) is one of the most direct and effective techniques for observing AGWs in the middle and upper atmosphere (Taylor and Hapgood, 1988; Gardner and Taylor, 1998). Since the first gravity wave imaging by Peterson and Adams (1983), ASAIs have become a key tool for studying gravity wave phenomena near the mesopause (Garcia et al., 1997; Coble et al., 1998; Tu C et al., 2009). Projects such as the Chinese Meridian Project have successfully established airglow observation

---

First author: Y. B. He, 2011536262@qq.com
Correspondence to: Q. C. Xu, xqc@nssc.ac.cn
Received 09 DEC 2025; Accepted 11 MAR 2026.

networks (Wang C et al., 2021, 2024), accumulating a substantial volume of valuable data and providing a solid foundation for the statistical study of gravity waves (Xu JY et al., 2015). This network has been instrumental in capturing and analyzing atmospheric waves from extreme events, providing key insights into the vertical coupling processes from the troposphere to the thermosphere. Notable examples include the study of multigroup atmospheric waves generated by the 2022 Hunga Tonga–Hunga Ha’apai volcanic eruption (Li QZ et al., 2024) and the investigation of gravity wave propagation from a typhoon to the upper atmosphere (Li QZ et al., 2022).

With the expansion of ground-based airglow imager networks, an accurate and efficient rapid-screening method to identify gravity wave events from the massive volume of airglow images is urgently needed. Traditional analysis methods rely on manual preliminary screening followed by least-squares harmonic fitting (Yamashita et al., 2013; Matsuda et al., 2017) or two-dimensional (2D; Krebsbach and Preusse, 2007; Schroeder et al., 2009) and three-dimensional (3D) Fourier analysis (John and Kumar, 2012; Matsuda et al., 2017), which is inefficient and highly subjective (Li QZ et al., 2025). With advances in deep learning, convolutional neural networks (CNNs) and object detection models have been utilized for the automatic identification of gravity waves in airglow images (González et al., 2022a). Sathish (2022) utilized a CNN for classifying AGWs from space-borne Synthetic Aperture Radar imagery, but the results were suboptimal because of insufficient training data. González et al. (2022b) used transfer learning, comparing various models and training strategies, and found that a combination of the Inception-v3 CNN architecture and a learning rate scheduler yielded the best performance, achieving a test accuracy of 75.85%. Xiao BM et al. (2024) further compared transfer learning with the YOLOv10 object detection model and concluded that transfer learning outperforms object detection methods under data-limited conditions. For ground-based observations, Lai C et al. (2019) developed an automated extraction procedure based on CNNs and Faster R-CNN (region-based CNN). By implementing uniform identification criteria, they mitigated the random errors associated with manual interpretation. However, because relatively strict thresholds were applied, the number of gravity waves identified was only 28.9% of the manual count. Chen JS et al. (2023) proposed an identification method combining improved CycleGAN (Cycle-Consistent Adversarial Network) data augmentation and the YOLOv5s object detector, achieving a mean average precision of 75.8%, with both detection speed and accuracy surpassing other mainstream object detection algorithms. Although deep learning demonstrates considerable potential for the automatic identification of gravity waves in ground-based airglow images, current methods encounter two core bottlenecks: First, the scarcity of annotated data makes complex models prone to overfitting, leading to poor generalization performance. Second, in the pursuit of high accuracy, models are often designed with excessive complexity, making it difficult to meet the stringent real-time processing requirements of ground-based stations. Furthermore, the ground-based observation environment is more complicated because of substantial stellar noise and weather interference. Current preprocessing methods lack specificity for these challenges, which increase the difficulty of effective feature extraction.

To address the aforementioned challenges, an efficient and reliable real-time identification framework for gravity waves in ground-based airglow images was proposed, representing the first systematic application of a transfer learning strategy to ground-based airglow images. By performing layered fine-tuning on an EfficientNetB3 backbone network pretrained on ImageNet (a publicly available image database), we effectively leveraged the prior knowledge from large-scale datasets, significantly mitigating the overfitting problem common with small sample sizes. In terms of model design, we eschewed computationally complex general-purpose networks (e.g., Inception-v3) in favor of the more lightweight EfficientNetB3, incorporating deformable convolution modules to maintain high sensitivity to gravity wave features while ensuring that computational efficiency would meet the real-time requirements for edge deployment. To address the low signal-to-noise ratio (SNR) in ground-based images, we proposed a wavelength-adaptive median filtering preprocessing method, which dynamically optimizes the denoising process based on the physical scale of the gravity waves. The final framework achieved an accuracy of 91.2% on the test set, reaching a millisecond-level inference speed. It has been successfully integrated into the monitoring network, enabling automatic, continuous, and real-time identification of gravity wave events.

2. Materials and Methods

2.1 Instrument and Dataset

The airglow images used in this study originated from ground-based ASAIs in the Lhasa area. The observation bands cover typical gravity wave-sensitive emission layers, such as OI 630 nm and OH (near-infrared). The station cameras are equipped with fish-eye lenses to obtain two-dimensional radiance distributions covering nearly the entire sky dome, with a single-frame field of view close to 180°. The typical exposure time is 60 s, with a temporal resolution of approximately 0.5–2 min. A single station can obtain hundreds to thousands of images per night, providing a sufficient temporal sampling density for resolving gravity wave morphology, propagation direction, and phase velocity in the middle and upper atmosphere (80–300 km; Narayanan and Gurubaran, 2013).

Raw data first underwent observational quality control (including star sensitivity drift calibration, cloud cover threshold determination, lunar phase screening, and elimination of nights with severe light pollution). Images from continuous time windows were selected for differential processing. Semiautomatic labeling was then performed based on the presence or absence of distinct banded, vortex, or ripple-type airglow perturbations, forming a binary-labeled image dataset (682 images with waves, 2561 images without waves). To avoid bias caused by severe class imbalance, negative samples were subjected to random undersampling and time-decorrelation processing during dataset construction, ensuring that the class ratios in the training, validation, and test sets remained roughly balanced overall. We utilized stratified sampling to divide the data into training (70%), validation (20%), and independent test (10%) sets, ensuring that samples from different nights would be distributed across different subsets as much as possible to mitigate the impact of temporal autocorrelation on the evaluation results.

2.2 Deep Learning Model

Transfer learning played a key role in improving the accuracy and robustness of automatic identification of gravity waves in ground-based airglow images. EfficientNet-B3 was selected as the backbone network. Its core advantage lies in its compound scaling strategy, which performs balanced scaling across three dimensions: network depth, width, and input resolution (Tan MX and Le, 2020). The MBConv (Mobile Inverted Bottleneck Convolution) blocks in EfficientNet also keep the inverted residual structure and depthwise separable convolution from the CNN MobileNetV2, which achieve significant compression of parameters and FLOPs (Floating Point Operations). Additionally, they incorporate Squeeze-and-Excitation channel attention blocks, enhancing the ability of the network to adaptively model the importance of features across channels (Sandler et al., 2019). Compared with classical CNN architectures such as ResNet-50, VGG, and Inception-v3, EfficientNet-B3 achieves a superior accuracy–efficiency trade-off with fewer parameters and a lower computational cost. This characteristic is crucial for deploying models on resource-constrained field stations or for near-real-time operational processing. Extensive experiments have shown that such lightweight, efficient networks with channel attention possess good transfer learning and generalization capabilities in optical remote sensing, medical imaging, and low SNR visual tasks (Qin Y, 2023). They are thus well-suited for AGWs, which are targets rich in textural details, spanning large scales, and sensitive to contrast variations.

To enhance the model's perception of wavelike and nonrigid structures, a deformable convolution module was utilized after the high-level semantic feature maps output by the backbone network. This module incorporates an additional DeformableConv 2D layer (deformable convolutional layer with 2D offsets), enabling dynamic adjustment of spatial sampling locations, thereby granting the network the ability to adaptively model complex wave patterns. The core of the DeformableConv2D layer involves learning a set of trainable offsets that shift the sampling locations of the standard convolution kernel, allowing the convolution to perceive and align with local geometric variations in the input feature maps (Dai JF et al., 2017). This mechanism is particularly suitable for handling scenarios lacking strict translation invariance, such as bending wavefronts, converging or diverging wave beams, and areas with strong background gradients. Through flexible control of the receptive field's sampling points, the model can more accurately analyze and reconstruct nonrigid patterns. Furthermore, to ensure effective convergence of the deformable module, we introduced additional L2 regularization constraints and an offset clipping mechanism to further limit the range of spatial deformation and prevent uncontrolled offsets from adversely affecting high-level semantic structures.

3. Data Processing and Model Training

3.1 Data Preprocessing

To address the challenges of a low SNR and strong stellar interference in ground-based airglow images, we utilized a targeted image enhancement pipeline prior to model input. This pipeline consisted of two main steps designed to highlight wavelike structures while suppressing background noise.

First, to mitigate random high-frequency noise and stellar point sources, we applied a multiscale adaptive median filtering technique. The filter window sizes were empirically selected as 9, 15, and 21 pixels. These scales correspond approximately to the typical horizontal wavelengths of gravity waves projected onto the image grid. By fusing the filtering results from these different scales, the method effectively preserves the morphological integrity of gravity wave bands while removing smaller stellar spots.

Second, to further isolate the transient gravity wave signals from the slowly varying background (such as the Milky Way, zodiacal light, and thin clouds), we utilized a dynamic frame differencing method. By subtracting the previous frame from the current frame within a continuous observation sequence, static or quasi-static background features were cancelled out. This process significantly enhanced the contrast of the propagating wave fronts, making the weak ripples more distinguishable for the subsequent neural network.

During model training, lightweight augmentation strategies were introduced to simulate observing conditions under different lunar phases and thin cloud cover. The data processing workflow is illustrated in Figure 1, which demonstrates the effects of different augmentation operations (flip, brightness, and contrast scaling) applied to a single representative source image.

3.2 Model Architecture

The model architecture, shown in Figure 2, primarily consists of three modules: feature extraction, deformable convolution, and classification decision. The input airglow images, after data augmentation operations, were uniformly resized to 300 × 300 pixels to match the input size of the pretrained EfficientNetB3. The extracted feature maps were then fed into a deformable convolution layer to further enhance the capability for gravity wave recognition. After global average pooling, the features were input to the classification head, which finally output a discrimination confidence score. A score above 0.5 was classified as "no wave," and otherwise as a gravity wave event occurrence.

3.3 Training Strategy

Considering the requirements for real-time deployment at stations and the risk of overfitting, we adopted a lightweight binary classification head. The custom Swish activation function is defined as $\text{Swish}\,(x) = x \cdot \sigma\,(\beta x)$, where σ is the sigmoid function, x denotes the input to the activation function, and the hyperparameter β controls the shape of the function. The L2 regularization penalty term was applied to the weight parameters θ_i of the fully connected layer as $\mathcal{L}_{\text{reg}} = \lambda \sum_i \theta_i^2$, where θ_i represents each individual weight parameter in the fully connected layer, and the hyperparameter λ controls the strength of the L2 penalty. Compared with traditional stacks of multilayers, high-dimensional fully connected layers significantly reduce the number of parame-

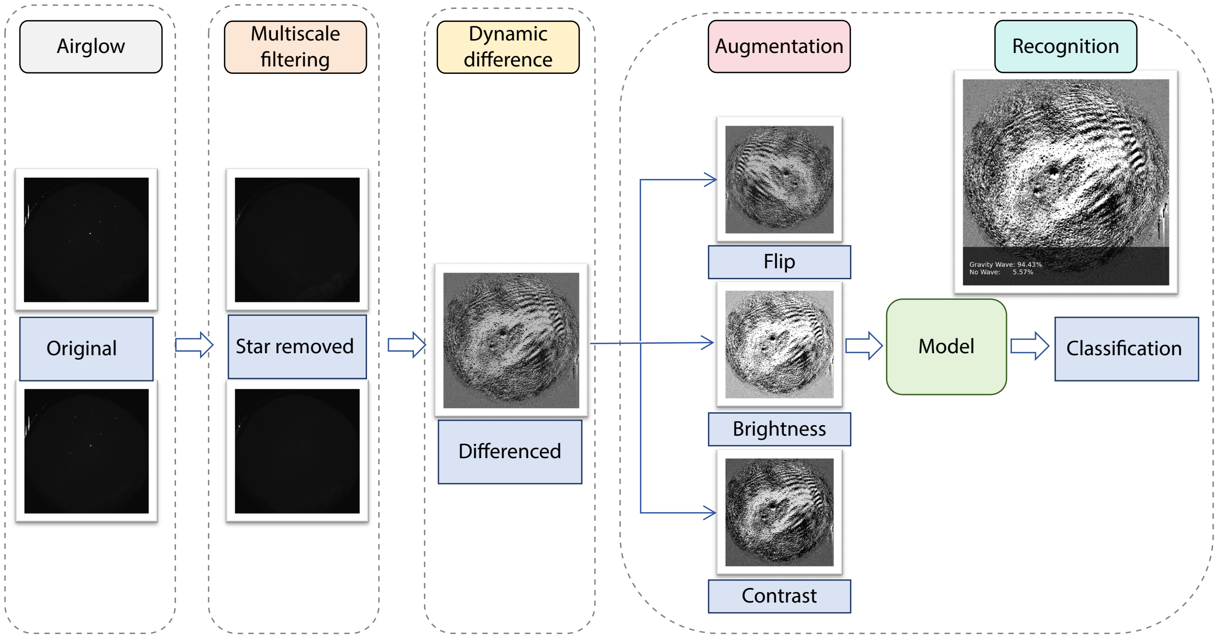

Figure 1. Flowchart of the rapid identification process for gravity waves in airglow images.

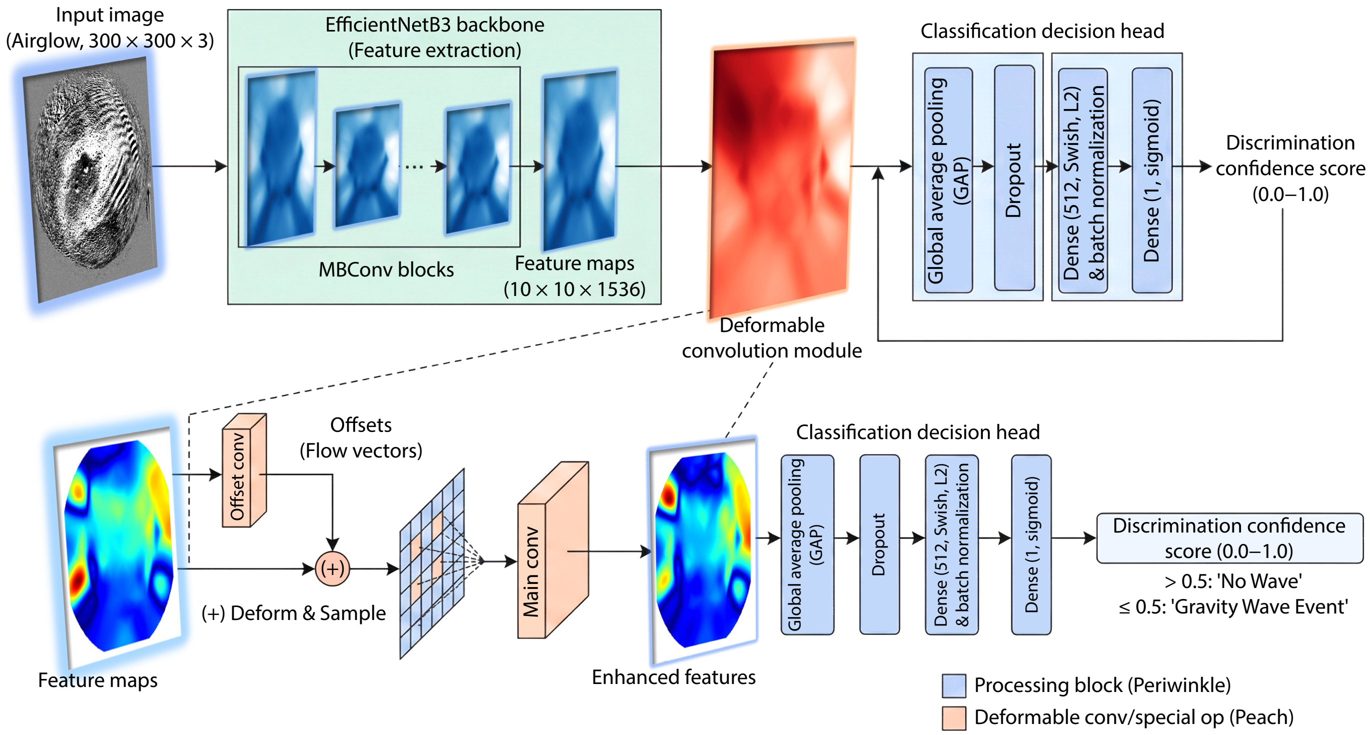

Figure 2. Model structure based on EfficientNetB3 and the deformable convolution module.

ters and GPU memory usage, and they typically yield more stable probability calibration in binary classification tasks. The dropout rate and weight decay are adjusted hierarchically at different stages to maximally suppress overfitting while preserving representational capacity.

To address the issues of scarce labeled samples and weak signal features, we designed a two-stage transfer fine-tuning method:

(1) *Initial stage (frozen backbone)*: Initially, all convolutional layers of EfficientNetB3 are frozen, and only the newly added lightweight classification head is trained. This stage uses a relatively high initial learning rate combined with the AdamW optimizer for rapid convergence, aiming to allow the new, fully connected layers to first adapt to the data distribution and establish basic decision boundaries.

(2) *Fine-tuning stage (partial unfreezing)*: After the classification head is trained, the higher level feature extraction modules at the

back end of EfficientNetB3 (e.g., layers after stage 200) are unfrozen. This keeps the low-level edge or texture and other features stable while allowing the high-level features to semantically adapt to the morphological features of AGWs (bands, coherent ripples, arc structures, etc.). This stage uses a lower learning rate, along with gradient clipping and layer-wise weight decay to stabilize training. The custom DeformableConv2D layer, added to the high-level feature maps output by the backbone, is also unfrozen and fine-tuned only in this stage with stricter regularization and an even lower learning rate for a short number of epochs. This layer learns offsets for spatial sampling points, enabling the convolution kernel to adaptively align with nonrigid, non-translation-equivariant wave structures (especially in scenarios with curved wavefronts, wave beam convergence or divergence, and strong background gradients).

In the frozen stage, we implemented the Keras version of the AdamW optimizer $\theta_{t+1} = \theta_t - \eta(\nabla L + \lambda\theta_t)$, which uses decoupled weight decay to help achieve better generalization on small- or medium-sized datasets and improves gradient stability to prevent overfitting (Loshchilov and Hutter, 2019). In the fine-tuning stage, a cosine annealing learning rate scheduler

$$\eta_t = \eta_{\min} + 0.5(\eta_{\max} - \eta_{\min})(1 + \cos(\pi t/T))$$

was selected to accelerate convergence, enhance generalization, and help escape local minima. A layer-wise weight decay mechanism was applied during fine-tuning, implementing an exponentially decaying strategy for the unfrozen layers $\lambda_l = 0.01 \cdot \mathrm{e}^{-2l/L}$. This process applied stronger decay to shallower layers (constraining overfitting on low-level features) and weaker decay to deeper layers (preserving high-level semantic features). Here, η is the learning rate controlling the parameter update step size, λ is the weight decay coefficient, T is the learning rate restart period, $\eta_{\max}$ and $\eta_{\min}$ are the maximum and minimum learning rates, l is the layer index (frozen layers are excluded from decay), and L is the total number of fine-tuned layers.

4. Results

The two-stage training strategy demonstrated stable convergence characteristics. In Stage 1, utilizing a cosine annealing scheduler, the validation loss decreased rapidly and plateaued after 30 epochs. In Stage 2, after unfreezing and fine-tuning the higher level features, the validation accuracy and F1-score showed further improvement. The introduction of the deformable convolution module enhanced the model's Recall on samples with complex backgrounds and weak signals, whereas the precision remained largely stable. This result indicates that the model improved the detection capability for challenging samples without excessively increasing false alarms. The training curves for loss and accuracy are shown in Figure 3, where EarlyStopping effectively prevented overfitting in the later stages.

The final performance of the model on the independent test set is summarized in Table 1.

To quantify the contribution of each design component to the final performance, a systematic ablation study was conducted. The results are illustrated in Figure 4. Using EfficientNetB3 as the baseline model, we achieved an accuracy of 82% without any specific optimizations. Training only the classification head with

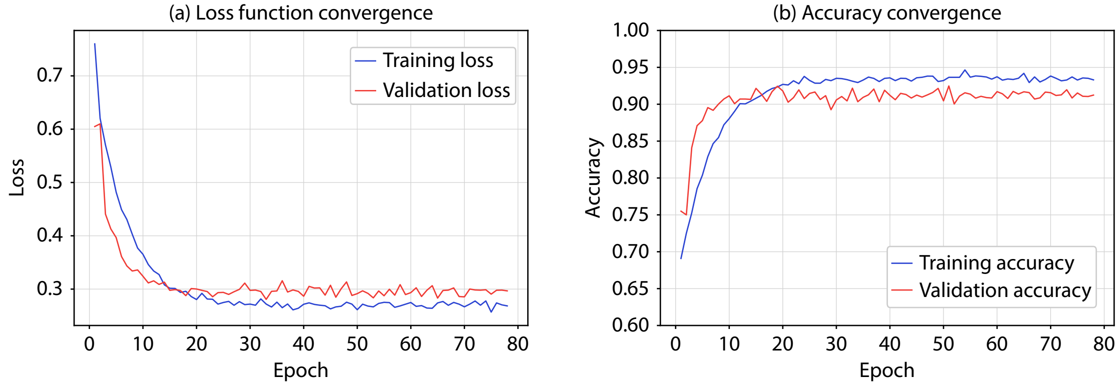
(a) Loss function convergence

(b) Accuracy convergence

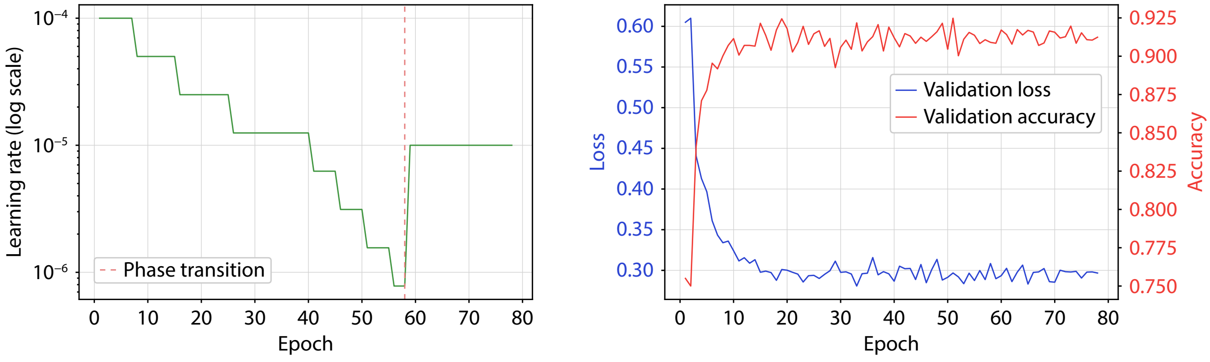
(c) Learning rate schedule

(d) Validation performance comparison

Figure 3. Training and validation loss and accuracy curves.

Table 1. Model performance metrics on the independent test set.

Metric	Value	95% confidence interval
Overall accuracy	91.24%	[89.1%, 93.4%]
F1-score	91.43%	[89.5%, 93.3%]
Wave detection rate (recall)	92.75%	[90.2%, 95.3%]
Wave detection precision	90.14%	[87.5%, 92.8%]
No-wave identification rate (specificity)	89.71%	[87.1%, 92.3%]
Validation loss	0.2967	[0.28, 0.32]

the feature extractor frozen yielded reasonable performance, but progressive learning achieved by unfreezing higher level features (Stage 2) provided additional gains. Unfreezing the entire network at once (nonlayered) was more prone to overfitting in this small-sample scenario. For the test set containing large-scale gravity waves, the model incorporating DeformableConv2D showed a significant increase in Recall, demonstrating its better adaptability to gravity wave structures. The AdamW optimizer demonstrated better generalization ability under the same regularization strength and greater stability with learning rate annealing, being less susceptible to performance drops on the validation set. The Swish activation function provided a slight but stable gain in the classification head, especially in the later stages of fine-tuning, and it helped reduce fluctuations in validation metrics. By integrating all the optimization strategies, the model achieved an accuracy of 91.24% and an F1-score of 91.43% compared to the baseline. The significant improvement across all metrics demonstrates the effectiveness of the multicomponent cooperative optimization.

To interpret the decision-making process of the proposed deep learning framework and verify that it focuses on physically relevant features, we utilized Gradient-weighted Class Activation Mapping (Grad-CAM). Grad-CAM serves as a visualization technique that highlights the specific regions of an input image contributing most significantly to the model's final classification, thereby explaining where the network is looking. The resulting attention heatmaps are presented in Figure 5. For samples containing gravity waves, the attention of the model was predominantly concentrated on the distinct wave band regions. In high-confidence predictions (probability >0.9), the heatmaps precisely delineated the wavelike structures, confirming that the model had learned to recognize the specific textural patterns of AGWs. Even in medium-confidence instances, the model demonstrated remarkable sensitivity to weak signals, correctly locating faint wave packets. Conversely, for "no-wave" samples, the attention was diffusely distributed over areas of background noise without focusing on specific structural features. Notably, in high-confidence negative predictions, the model successfully ignored potential confounders, such as cloud interference and smooth gradients caused by lunar illumination. This capability confirms that the model relies on genuine wave morphology rather than environmental artifacts, providing strong support for the physical plausibility of the model's decisions and offering a reliable visual basis for subsequent physical inversions (e.g., phase velocity and propagation direction estimation).

The confusion matrix evaluated on the independent balanced test set is shown in Figure 6a. This test set represents a random 10% split derived from the balanced dataset (consisting of 682 positive samples and 682 undersampled negative instances), ensuring that the evaluation metrics reflected the performance of the model on an unbiased distribution. The results revealed the superior classification performance of the model: Among the positive instances, 92.8% of gravity waves were correctly identified (True Positives), and among the negative instances, 89.7% of no-wave samples were correctly identified (True Negatives). In the receiver operating characteristic (ROC) curve shown in Figure 6b, the area under the curve (AUC) value reached 0.964, indicating excellent classification capability. The proximity of the curve to the ideal top-left corner signifies robust model performance across varying discrimination thresholds.

As benefits from the efficient architecture of EfficientNetB3 and the lightweight classification head, the complete inference time for a single image in this study model was 23 ms. This value fully meets the requirements for practical operational applications and offers significant advantages for large-scale batch processing

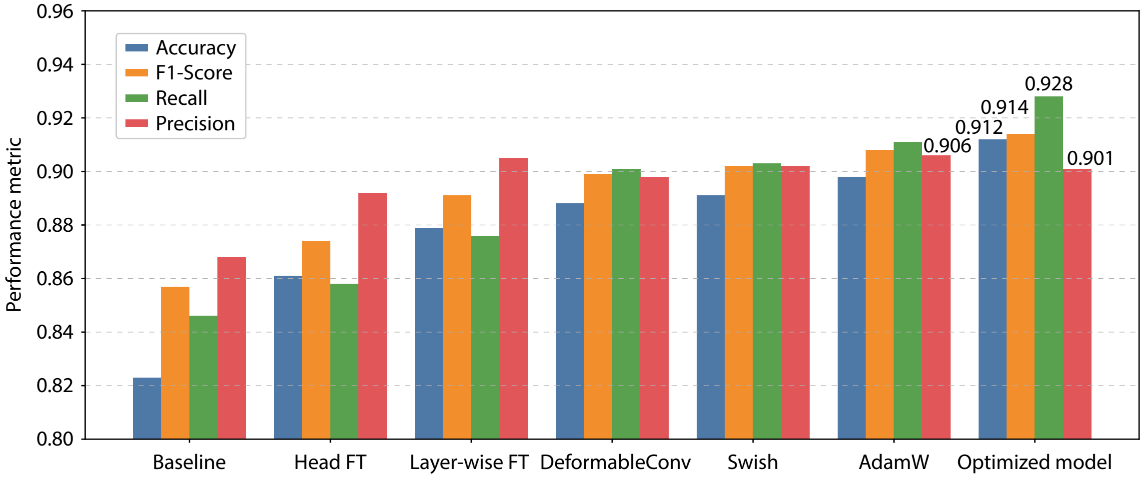

Figure 4. Comparative results of the ablation study.

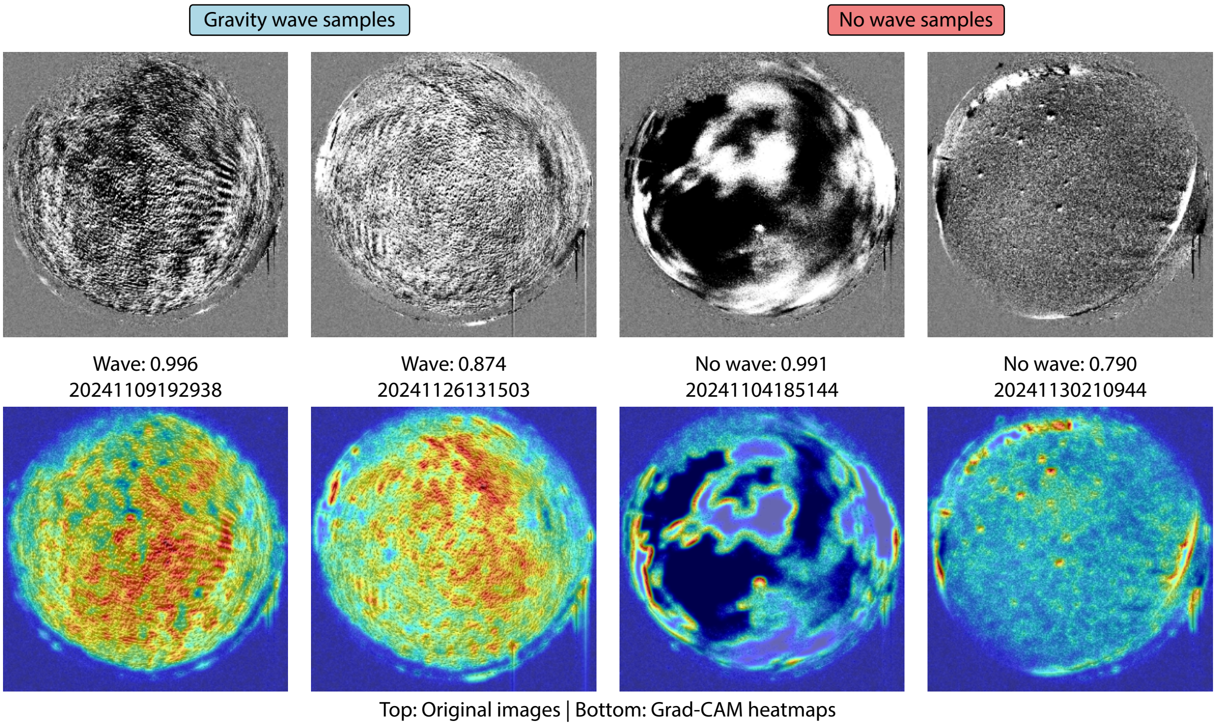

Figure 5. Grad-CAM heatmaps visualizing model attention.

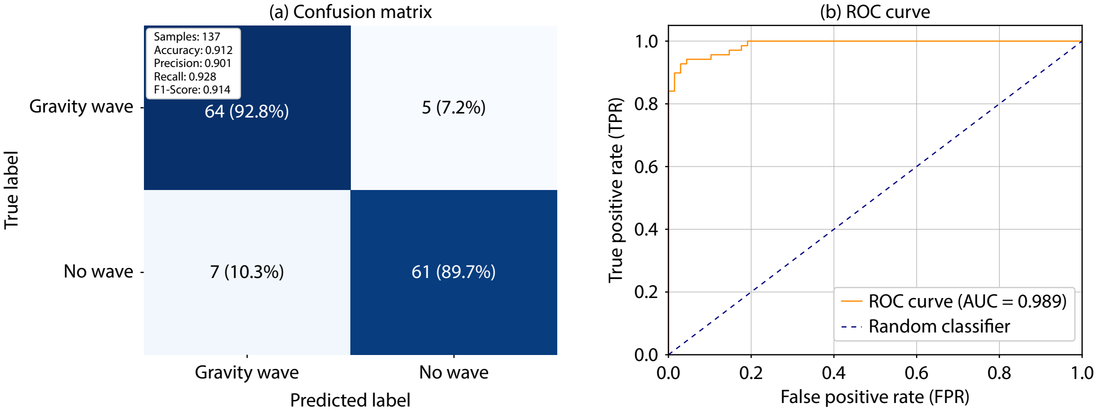

Figure 6. (a) Confusion matrix and (b) receiver operating characteristic (ROC) curve.

scenarios involving continuous hourly data throughout the night.

5. Conclusions

A real-time identification framework for ground-based airglow gravity waves was proposed. Integrating physics-guided preprocessing and lightweight deep learning, it effectively addressed the limitations of traditional methods in terms of efficiency and generalization ability. The core innovations can be summarized as follows:

(1) *Domain-adaptive preprocessing enhancement*: A multiscale fusion method combining wavelength-adaptive median filtering and dynamic differencing is utilized to address the low SNR challenge. Tailored to the typical wavelength range of gravity waves, this method suppresses noise, significantly enhancing the SNR of weak gravity wave signals while effectively removing stellar noise.

(2) *Lightweight model and hierarchical transfer learning strategy*: With the EfficientNetB3 backbone network and a custom deformable convolution module, coupled with a two-stage hierarchical fine-tuning mechanism, the framework enhances the perception and modeling capability for the nonrigid characteristics of gravity waves while effectively mitigating overfitting on small datasets.

(3) *High accuracy and real-time performance*: The accuracy rate of the constructed model on the independent test set reached 91.2%, and the reasoning time of a single image reached 23 ms, meeting the requirements of real-time operation processing.

Although the proposed framework demonstrated robust performance in real-time screening, we acknowledge certain limitations in the current study. First, the model currently relies on single-frame spatial features, potentially underutilizing the rich dynamic information embedded in the temporal evolution of gravity waves. Second, the binary classification head, although efficient, restricts the system to presence detection rather than fine-grained typification or parameter extraction. Third, like most supervised learning approaches, the performance of our model depends heavily on the quantity and quality of manual annotations.

To build on these observations, future work will focus on advancing the capabilities of the model in three key directions. First, to move beyond simple detection, we aim to develop integrated parameter inversion algorithms that can simultaneously perform fine-grained classification and estimate key physical parameters (e.g., phase velocity and propagation direction), thereby facilitating more in-depth climatological analysis. Second, spatiotemporal modeling techniques, such as 3D-CNNs or temporal Transformers, will be incorporated to leverage sequential information for better noise suppression and motion characterization. Additionally, we plan to explore self-supervised learning strategies to reduce the dependence on manual labeling and improve generalization across diverse observational conditions. The ultimate goal is to construct a fully automated, physics-informed analysis platform that seamlessly integrates detection, identification, and quantitative parameter extraction.

Acknowledgments

This work was supported by the Strategic Priority Research Program of the Chinese Academy of Sciences (Grant No. XDA17010302) and the National Natural Science Foundation of China (Grant Nos. 12241101, 42174192, and 11872128).

References

Ayorinde, T. T., Wrasse, C. M., Takahashi, H., Barros, D., Figueiredo, C. A. O. B., Da Silva, L. A., and Bilibio, A. V. (2024). Investigation of the long-term variation of gravity waves over South America using empirical orthogonal function analysis. *Earth Planets Space*, *76*(1), 105. https://doi.org/10.1186/s40623-024-02045-0

Chen, J. S., Ma, W. Z., Fang, S. F., and Zou, Z. M. (2023). Object recognition of atmospheric gravity wave based on foundation airglow images. *Comput. Eng. (in Chinese)*, *49*(11), 13–23. https://doi.org/10.19678/j.issn.1000-3428.0066521

Coble, M., Papen, G. C., and Gardner, C. S. (1998). Computing two-dimensional unambiguous horizontal wavenumber spectra from OH airglow images. *IEEE Trans. Geosci. Remote Sens.*, *36*(2), 368–382. https://doi.org/10.1109/36.662723

Dai, J. F., Qi, H. Z., Xiong, Y. W., Li, Y., Zhang, G. D., Hu, H., and Wei, Y. C. (2017). Deformable convolutional networks. In *Proceedings of the 2017 IEEE International Conference on Computer Vision* (pp. 764–773). Venice, Italy: IEEE. https://doi.org/10.1109/ICCV.2017.89

Duan, B. Q., and Pan, W. L. (2017). A data preprocessing method and preliminary results of all-sky airglow image. *Chin. J. Space Sci. (in Chinese)*, *37*(1), 94–104. https://doi.org/10.11728/cjss2017.01.094

Fritts, D. C. (1984). Gravity wave saturation in the middle atmosphere: A review of theory and observations. *Rev. Geophys.*, *22*(3), 275–308. https://doi.org/10.1029/RG022i003p00275

Fritts, D. C., and Alexander, M. J. (2003). Gravity wave dynamics and effects in the middle atmosphere. *Rev. Geophys.*, *41*(1), 1003. https://doi.org/10.1029/2001RG000106

Garcia, F. J., Taylor, M. J., and Kelley, M. C. (1997). Two-dimensional spectral analysis of mesospheric airglow image data. *Appl. Opt.*, *36*(29), 7374–7385. https://doi.org/10.1364/AO.36.007374

Gardner, C. S., and Taylor, M. J. (1998). Observational limits for lidar, radar, and airglow imager measurements of gravity wave parameters. *J. Geophys. Res.: Atmos.*, *103*(D6), 6427–6437. https://doi.org/10.1029/97JD03378

González, J.L., Chapman, T., Chen, K., Nguyen, H., Chambers, L., Mostafa, S.A.M., Wang, J., Purushotham, S., Wang, C., and Yue, J. (2022a). Exploring machine learning based atmospheric gravity wave detection. https://doi.org/10.13016/M2OZLY-WDQU

González, J. L., Chapman, T., Chen, K., Nguyen, H., Chambers, L., Mostafa, S. A. M., Wang, J. W., Purushotham, S., Wang, C. X., and Yue, J. (2022b). Atmospheric gravity wave detection using transfer learning techniques. In *2022 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)* (pp. 128–137). Vancouver, WA, USA: IEEE. https://doi.org/10.1109/BDCAT56447.2022.00023

John, S. R., and Kumar, K. K. (2012). TIMED/SABER observations of global gravity wave climatology and their interannual variability from stratosphere to mesosphere lower thermosphere. *Climate Dyn.*, *39*(6), 1489–1505. https://doi.org/10.1007/s00382-012-1329-9

Krebsbach, M., and Preusse, P. (2007). Spectral analysis of gravity wave activity in SABER temperature data. *Geophys. Res. Lett.*, *34*(3), L03814. https://doi.org/10.1029/2006GL028040

Lai, C., Xu, J. Y., Yue, J., Yuan, W., Liu, X., Li, W., and Li, Q. Z. (2019). Automatic extraction of gravity waves from all-sky airglow image based on machine learning. *Remote Sens*, *11*(13), 1516. https://doi.org/10.3390/rs11131516

Li, Q. Z., Xu, J. Y., Liu, H. L., Liu, X., and Yuan, W. (2022). How do gravity waves triggered by a typhoon propagate from the troposphere to the upper atmosphere?. *Atmos. Chem. Phys.*, *22*(18), 12077–12091. https://doi.org/10.5194/acp-22-12077-2022

Li, Q. Z., Xu, J. Y., Gusman, A. R., Liu, H. L., Yuan, W., Liu, W. J., Zhu, Y. J., and Liu, X. (2024). Upper-atmosphere responses to the 2022 Hunga Tonga–Hunga Ha'apai volcanic eruption via acoustic gravity waves and air–sea interaction. *Atmos. Chem. Phys.*, *24*(14), 8343–8361. https://doi.org/10.5194/acp-24-8343-2024

Li, Q. Z., Xu, J. Y., Yuan, W., Liu, X., Zhu, Y. J., and Liu, W. J. (2025). Three-dimensional spectral analysis of gravity waves from airglow observations over Northwest China. *Earth Planet. Phys.*, *9*(4), 988–994. https://doi.org/10.26464/epp2025050

Loshchilov, I., and Hutter, F. (2019). Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101. https://doi.org/10.48550/arXiv.1711.05101

Matsuda, T. S., Nakamura, T., Ejiri, M. K., Tsutsumi, M., Tomikawa, Y., Taylor, M. J., Zhao, Y. C., Pautet, P. D., Murphy, D. J., and Moffat-Griffin, T. (2017). Characteristics of mesospheric gravity waves over Antarctica observed by Antarctic Gravity Wave Instrument Network imagers using 3-D spectral analyses. *J. Geophys. Res.: Atmos.*, *122*(17), 8969–8981. https://doi.org/10.1002/2016JD026217

Nappo, C. J. (2002). *An Introduction to Atmospheric Gravity Waves*. Amsterdam, Netherlands: Academic Press.

Narayanan, V. L., and Gurubaran, S. (2013). Statistical characteristics of high frequency gravity waves observed by OH airglow imaging from Tirunelveli (8.7°N). *J. Atmos. Sol. Terr. Phys.*, *92*, 43–50. https://doi.org/10.1016/j.jastp.2012.09.002

Peterson, A. W., and Adams, G. W. (1983). OH airglow phenomena during the 5–6 July 1982 total lunar eclipse. *Appl. Opt.*, *22*(17), 2682–2685. https://doi.org/10.1364/AO.22.002682

Qin, Y. (2023). Research progress of deep learning in liver semantic segmentation and volume measurement based on medical image. *Adv. Clin. Med. (in Chinese)*, *13*(7), 11581–11587. https://doi.org/10.12677/ACM.2023.1371620

Sandler, M., Howard, A., Zhu, M. L., Zhmoginov, A., and Chen, L. C. (2019). MobileNetV2: Inverted residuals and linear bottlenecks. arXiv preprint arXiv:1801.04381. https://doi.org/10.48550/ARXIV.1801.04381

Sathish, A. K. (2022). *Detection of Atmospheric Gravity Waves*. Delft, Netherlands: Delft University of Technology.

Schroeder, S., Preusse, P., Ern, M., and Riese, M. (2009). Gravity waves resolved in ECMWF and measured by SABER. *Geophys. Res. Lett.*, *36*(10), L10805.

https://doi.org/10.1029/2008GL037054

Tan, M. X., and Le, Q. V. (2020). EfficientNet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946. https://doi.org/10.48550/ARXIV.1905.11946

Taylor, M. J., and Hapgood, M. A. (1988). Identification of a thunderstorm as a source of short period gravity waves in the upper atmospheric nightglow emissions. *Planet. Space Sci.*, *36*(10), 975–985. https://doi.org/10.1016/0032-0633(88)90035-9

Tu, C., Hu, X., Guo, S. Y., Yan, Z. A., and Cheng, Y. Q. (2009). CSSAR airglow gravity wave imager and its preliminary observation. In *Proceedings Volume 7384, International Symposium on Photoelectronic Detection and Imaging 2009* (pp. 421–428). Beijing, China: SPIE. https://doi.org/10.1117/12.835580

Tu, C., and Hu, X. (2012). Spectra of the OH airglow perturbation on 5 January 2009 at Langfang. *Chin. J. Space Sci. (in Chinese)*, *32*(6), 824–828. https://doi.org/10.11728/cjss2012.06.824

Wang, C., Chen, Z. Q., Hu, L. H., Hu, Z. J., Xue, X. H., Xu, Q. C., Zhang, X. G., Li, F. Q., Wang, W., … Gong, W. L. (2021). Development and prospect of China's space-based and ground-based space environment monitoring platforms. *Spacecr. Environ. Eng. (in Chinese)*, *38*(3), 225–239. https://doi.org/10.12126/see.2021.03.001

Wang, C., Xu, J. Y., Chen, Z. Q., Li, H., Feng, X. S., Huang, Z. H., and Wang, J. Y. (2024). China's ground-based space environment monitoring network—Chinese Meridian Project (CMP). *Space Weather*, *22*(7), e2024SW003972. https://doi.org/10.1029/2024SW003972

Xiao, B. M., Hu, S. S., Ai, W. H., and Li, Y. (2024). Atmospheric gravity wave detection in low-light images: A transfer learning approach. *Electronics*, *13*(20), 4030. https://doi.org/10.3390/electronics13204030

Xu, J. Y., Li, Q. Z., Yue, J., Hoffmann, L., Straka, W. C., III, Wang, C. M., Liu, M. H., Yuan, W., Han, S., … Ning, B. Q. (2015). Concentric gravity waves over northern China observed by an airglow imager network and satellites. *J. Geophys. Res.: Atmos.*, *120*(21), 11058–11078. https://doi.org/10.1002/2015JD023786

Yamashita, C., England, S. L., Immel, T. J., and Chang, L. C. (2013). Gravity wave variations during elevated stratopause events using SABER observations. *J. Geophys. Res.: Atmos.*, *118*(11), 5287–5303. https://doi.org/10.1002/jgrd.50474