

國立臺灣大學理學院大氣科學系

博士論文

Department of Atmospheric Sciences

College of Science

National Taiwan University

Doctoral Dissertation



應用可解釋深度學習

識別大氣過程中的非線性形態特徵

Explainable Deep Learning for Identifying

Nonlinear Morphological Features of Atmospheric Processes

陳逸昌

Chen, Yi-Chang

指導教授：吳健銘 博士

Advisor: Wu, Chien-Ming, Ph.D.

中華民國 115 年 01 月

January, 2026

致謝



當年踏入大氣科學系時，我從未想過這段求學之路會一路延伸至博士學位。這不只是一段學術訓練的過程，更是一段因理解、幫助與鼓勵而得以延續的旅程。

旅程的起點，是一位仍在摸索方向的學士班學生，帶著忐忑與期待，敲開新進教師的辦公室。感謝吳健銘老師願意接納當時還懵懂的我。即便那時的我對大氣研究幾乎一無所知，老師仍抱持著「不妨試試看」的信念，給了我一個開始的機會，並擔任了我的指導教授。從學士專題到碩士班，歷經服役、留學、工作，再到博士班的每一個階段，這份「試試看」的信念始終伴隨著我。尤其在研究、工作或生活遭遇低潮、迷惘與自我懷疑的時候，我總會想到去找老師聊一聊，很多時候，只要坐下來把話說出口，那些卡住的地方就會慢慢鬆開。老師的支持與引導，以及時不時的靈魂拷問，讓我在面對這些關卡時仍能保有當年嘗試的初心，學會在不確定之中繼續前行。

旅程需要有方向，師長們的提醒與推動，指導與鼓勵，是引導我走向正確的道路的關鍵。感謝郭鴻基老師以宏觀而深遠的視角提供建議，讓我能站在我無法企及的高度重新審視我的研究方向與意義。感謝陳維婷老師在博士班歷程中同時給予嚴謹的科學建議與溫暖的關懷，推動我發表第一篇論文，並給予我一個擔任課程助教的機會，讓我能對自己未來的生涯規劃進行更具體的探索。感謝蘇世顥老師在博士班的求學階段與未來發展上給予許多指點，雖然一年間會見面的次數不多，但每一次的談話都讓我在重要的轉折處能更實際地思考研究的方向，以及未來的生涯規劃。感謝梁禹喬老師邀請我投入不同領域的研究工作，跨足大尺度及平流層的研究，使我的旅程中多了一些不同的風景。

C302.5 研究室的夥伴們，也是我這趟旅途中不可或缺的重要存在。謝謝每一位在討論中給予我意見、在困難時伸出援手、在日常研究生活中陪我一起努力的人。感謝俊彥總能溫和而不失專業地提供學術的建議，以及最後擔任我的口試委員；感謝曼耕分享許多職場人生經驗，讓我在職涯發展迷惘之際，能從過來人的經驗獲得啟發；感謝柏言與我互訴同時身處學術與職場的甘苦，使我在忙碌之餘仍能得到理解與共鳴；感謝金德協助處理資料，分擔了研究中許多困難而關鍵的工作；感謝少禹、鵬任、冠廷、巧薇、懷逸、毓琇、定融、佳欣、佑軒、富

聖、宇泓、清翔、峻愷、彥丞、子強、宣諭、晉維、宗育、志廷、適文、維傑，以及所有參與我研究中討論的人們。因為你們的存在，讓研究室不只是工作的場所，有人願意聽我嘮叨碎念、有人願意陪我一起想辦法，也有人提醒我不要把自己逼得太緊，讓我得以喘息、重新出發的地方。



在這段旅程中，邊工作邊攻讀博士班，必然會有無數個「蠟燭兩頭燒」的情況。感謝博士班期間任職的兩間公司同仁與主管們的包容與體諒，讓我在工作與學業之間還能維持喘息的空間。同時，對於不熟悉學校行政程序的我，在行政作業上也少不了美雲姊、國豐助教、明昌助教的協助，讓許多繁瑣的流程得以順利推進。

也謝謝一路陪伴在我身旁的朋友們。感謝達陽、忠杰、秩群、威鎮、與我共度許多珍貴的休閒時光；在人生中，能有認識這麼久並且仍能持續分享彼此生活的朋友並不容易，你們的存在讓我在高壓之下仍能記得生活的溫度。感謝家馨從博班入學考試到畢業之間一路的相伴，讓我在關鍵的每一步都不覺得孤單。也感謝在深夜與我語音聊天的推友們，那些挑燈夜戰、研究停滯的夜晚，能找到人說說話，沉澱思緒，往往就能重新獲得許多力量與前行的動力。

最後，我想把感謝獻給我的家人。感謝爸爸、媽媽、妹妹無條件支持我的選擇。博士班充滿了許多看不見終點的分岔路，每一次抉擇都可能牽動未來截然不同的方向，但家人的理解與信任，總能讓我在最需要力量的時候，重新找回面對的勇氣，而不是獨自一人徬徨地面對著十字路口。也感謝泡泡在我心情低落孤單的時候陪在我身邊，那些靜謐的凌晨時光也是旅途中的一番風景。

研究的旅程，從來都不是走在一條筆直平穩的道路上，往往是伴隨著反覆嘗試、長時間的停滯與挫折循環。那些看似沒有進展的日子裡，會開始懷疑自己的能力與選擇；然而，正因為在旅途中遇見了許多願意伸出援手的人們，願意一起走一段的人們，願意佇足等待我的人，我才得以欣賞到旅程中的壯麗風景，才能體驗到旅程中的精采時刻，才有辦法一次次重新站起，把困難化為前進的動力。

謹以此論文，獻給所有成就這趟旅程的人們。

摘要



近年來，深度學習技術在大氣科學領域研究中的應用迅速增長，然而，這些模型往往被視為缺乏物理透明度的「黑盒子」，限制了其在科學推論上的可靠性。本論文旨在建立一個可解釋的深度學習框架，提出以「形態學（Morphology）」作為連結數據驅動預測與大氣物理過程的關鍵指標。透過捕捉物理變數在空間上的非線性形態特徵，我們得以在不同尺度下揭示大氣動力過程的演變機制。

本研究針對三個不同的大氣尺度，分別採用了三種深度學習架構來驗證此框架的可行性。首先，在大尺度方面，針對突發性平流層增溫（SSW）現象，我們利用「卷積變分自動編碼器（Convolutional Variational Autoencoder）」構建了平流層極地渦旋的潛在變數空間相位圖。結果顯示，VAE 能有效捕捉渦旋從位移型到分裂型的非線性連續演變路徑，其表現優於傳統的線性主成分分析（PCA），為極地渦旋的形態演變研究，以及未來的預測研究提供了新的視角。

其次，在中尺度方面，針對熱帶對流聚合現象，我們開發了一套迭代式特徵移除的卷積神經網路（CNN）框架。透過反覆訓練與遮蔽輸入資料中的顯著特徵（如平均雲水含量），我們嘗試找出隱藏在神經網路決策背後的關鍵形態因子。研究發現，除了雲覆蓋率外，雲的邊緣複雜度（碎形維度）是模型辨識對流是否集結的一個重要的非線性特徵。

最後，在小尺度方面，針對深對流中的冷池（Cold Pool）動力過程，我們應用 U-Net 模型，嘗試解決全球風暴解析模式（GSRMs）中對於冷池的解析度不足的問題。本研究利用高解析度大渦模擬（LES）資料，成功從低解析度的環境場中重建出高解析度的冷池形態與強度分佈。同時，敏感度分析進一步顯示近地面上的動力場對於

冷池邊界的重建有一定的重要性，其影響力在此尺度下的冷池結構重建上大於熱力變數。



綜合以上，本論文展示了深度學習模型在經過適當設計後，不僅能作為預測工具，更能成為探索大氣非線性形態特徵的分析工具。此框架透過捕捉大尺度物理場演變過程、擷取中尺度對流特徵、到嘗試重建次網格細節，期望能為提升數值模式的參數化的未來發展與深度學習可解釋性提供貢獻。

關鍵詞：可解釋深度學習、大氣形態特徵、卷積神經網路、變分自動編碼器、冷池、突發性平流層增溫

Abstract



The application of deep learning in atmospheric sciences has expanded rapidly in recent years. However, these models are often treated as "black boxes" lacking physical transparency, which limits their reliability for scientific inference. This dissertation aims to establish an explainable deep learning framework by proposing "morphology" as a key physical indicator to bridge the gap between data-driven predictions and atmospheric physical processes.

This research validates the proposed framework through three distinct atmospheric scales using appropriate deep learning architectures. First, at the large scale, focusing on Sudden Stratospheric Warmings (SSWs), we use a Convolutional Variational Autoencoder (VAE) to construct a latent space phase diagram of the stratospheric polar vortex. The results demonstrate that the VAE effectively captures the nonlinear continuous evolution of the vortex—transitioning between displacement and splitting events—outperforming Principal Component Analysis (PCA) and providing a novel perspective on vortex dynamics.

Second, at the mesoscale, investigating tropical convective aggregation, we develop an iterative feature-removal framework using Convolutional Neural Networks (CNNs). By iteratively training the model and masking dominant features (such as average cloud water path) from the input data, we aim to identify and extract the hidden morphological factors driving the network's decisions. This study reveals that, beyond cloud coverage, the complexity of cloud edges (fractal dimension) serves as an important nonlinear feature for identifying the occurrence of the convective aggregation.

Finally, at the small scale, addressing cold pool dynamics within deep convection, we apply a U-Net model to deal with the resolution gap in Global Storm-Resolving Models (GSRMs). Using high-resolution Large-Eddy Simulation (LES) data, we successfully reconstruct high-resolution cold pool morphology and intensity distributions from coarse-grained environmental fields. Sensitivity analysis further reveals that the near-surface dynamic field (wind divergence signals) is important for reconstructing cold pool boundaries, playing a more critical role than thermodynamic variables alone.

In conclusion, this dissertation demonstrates that when properly designed, deep learning models can serve not only as prediction tools but also as powerful physical analysis instruments for exploring nonlinear atmospheric morphology. By representing large-scale evolution, extracting mesoscale features, and reconstructing subgrid details, this framework is expected to provide a foundation for improving parameterization schemes and enhancing physical interpretability of deep learning research in atmospheric sciences in the future.

Keywords: Explainable Deep Learning, Atmospheric Morphology, Convolutional Neural Network, Variational Autoencoder, Cold Pool, Sudden Stratospheric Warming

Contents



致謝

i

摘要

iii

Abstract

v

Contents

vii

List of Figures

ix

List of Tables

xiv

1 Introduction

1.1 The Rise of Deep Learning in Atmospheric Sciences	1
1.2 The Problem of Black Box: The Need for Interpretability and Explainability	2
1.3 Morphology as a Physical Indicator: Linking Visual Features to Domain Knowledge	2
1.4 Thesis Objectives and Framework	3

2 Representing the Morphological Evolution of Large-Scale Circulation using Variational Autoencoder

2.1 Introduction	8
2.2 Data and Methods	9
2.3 Results: The Morphological Phase Diagram	15
2.4 Discussion	23
2.5 Summary and Conclusion	25

3 Extracting Nonlinear Morphological Features of Mesoscale Convection via an Iterative Deep Learning Framework

3.1 Introduction	27
3.2 Methodology	28
3.3 Iterative Extraction of Morphological Features	32
3.4 Discussion: Decoding the Hidden Layer	35

3.5	Summary	37
4	Capturing Subgrid-Scale Cold Pool Morphology: A U-Net Based Morphological Reconstruction	
4.1	Introduction	38
4.2	Methodology	38
4.3	Results: Reconstructing Sub-grid Morphology	44
4.4	Explainability: Identifying Key Morphological Features	47
4.5	Summary and Discussion	49
5	General Discussion	52
5.1	Linking multi-scale phenomena through morphology	52
5.2	The Role of Neural Network Architectures in Morphological Analysis	52
5.3	The Choice of Loss Functions	56
5.4	Morphology in Time	57
5.5	From Black Box to Glass Box: Bridging Data and Physics	58
6	Conclusion and Future Perspectives	61
6.1	Summary of Major Findings	61
6.2	Contributions	62
6.3	Future Perspectives	62
References		65

List of Figures



2-1 Illustration of the bilinear interpolation method applied to a Z10 field of a split-type SSW event from WACCM6 simulations. (a) The longitude–latitude map of the Northern Hemisphere in an orthographic projection with 90° N at the center, with the region of color shadings indicating the domain to be interpolated. (b) The resultant 36×36 squared grid after using the bilinear interpolation method. The parallels are respectively at 80, 60, and 40 degrees from the innermost to the outermost. 11

2-2 A schematic of the VAE structure used in this study. The top row represents the combination of a decoder and an encoder, which form the VAE model. The middle row illustrates the details of the encoder, while the bottom row those of the decoder. The vertical solid line after each layer denotes the max-pooling/upsampling layers, which reduce or increase the dimension of pixeled images, whereas the cubes represent the convolutional layers with varying numbers of channels. The encoder reads in a 36×36 Z10 input image, compresses its information through the convolutional layers, and outputs two means (μ_0, μ_1) and two variances (σ_0^2, σ_1^2) that form the latent space. The decoder utilizes these sampled results based on Gaussian distribution and generates a 36×36 Z10 output through multiple convolutional layers. The mean squared error (MSE) between the original and generated images and the Kullback–Leibler divergence between the latent space and the Gaussian distribution are computed. These errors are backpropagated to train the encoder and decoder. The light blue color shading highlights the two means of the latent space, which is used to produce the phase diagram in this study. 13

2-3 This figure illustrates the corresponding Z10 field generated by the VAE decoder when given the varying values of z_0 and z_1 in the latent space. The green circle in each panel represents the circle of 60° N. 15

2-4 (a) The evolution track of a displacement-type SSW event from WACCM6 simulations mapped onto the VAE phase diagram. (b) As in (a), but for the track mapped onto the PCA phase diagram. (c) The WACCM6 Z10 spatial patterns during days $[-10, 10]$. (d) The VAE reconstructed Z10 fields 17

during days $[-10, 10]$. (e) As in (d), but for PCA reconstructed Z10 fields. The values in the subtitle parenthesis in (c) are the aspect ratio (r) and central latitude (ϕ_o) of this SSW event, whereas those in (c) and (d) denote the pattern correlation (R) and RMSE. The green circle in each panel of (c)–(e) represents the circle of 60° N.

2-5 The same as in Figure 2-4, but for a split-type SSW. 19

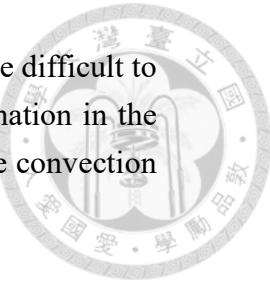
2-6 (a) The occurrence probability of the Z10 fields associated with all displacement-type SSWs from the WACCM6 testing dataset in the VAE phase diagram. The values are presented as percentages after normalized by the total number of events. (b) As in (a) but for split-type SSWs from WACCM6 simulations. Figures (c) and (d) are the same as (a) and (b), but for the PCA phase diagram. The blue dots are the Z10 fields from the reanalysis dataset mapped onto the VAE and PCA phase diagrams. 20

2-7 The pattern correlation (R) and root mean squared error (RMSE) for (a) and (b) the reconstructed displacement-type SSWs and (c) and (d) the reconstructed split-type SSWs from WACCM6 simulation against corresponding full Z10 during the period of days $[-10, 10]$. The color shading indicates the range of one standard deviation. 22

2-8 (a) Distributions of distance traveled in the VAE phase diagram for all split-type (magenta bars) and displacement-type (blue bars) SSWs. (b) Distributions of distance traveled for all simulated SSWs (blue bars) and observational (magenta bars) SSWs. (c) and (d) The same as (a) and (b) but for the distributions of distance traveled in the PCA phase diagram. 23

3-1 A schematic diagram shows the data flow of the CNN model and our framework. Each kernel of the convolution layers is 3×3 , and the activations are ReLUs. The pool size of the max-pooling layers is 3×3 , and the activation function of the output layer is sigmoid to represent the probability in $[0,1]$. The flowchart shows our framework for retrieving the physics process and the role of the CNN model in our framework. 28

3-2 (Non agg-1) to (non agg-4) shows the examples of the nonaggregated case, and (Agg-1) to (Agg-4) shows the aggregated case with the cloud water path values in $(\text{kg} \cdot \text{m}^{-2})$. Each subgraph shows the simulations' 512×512 30



(km) domain. The solid red lines represent the area that may be difficult to classify visually. Those areas can be the clouds before elimination in the nonaggregation cases, and the clouds are not at the core of the convection in the aggregation cases.

3-3 The model predicted probabilities of the aggregation (y-axis in each subgraph) from the CNN model in the different iterations and the potential property derived from the CWP values (x-axis in each subgraph). Blue points in each subgraph stand for the actual label of the training data are nonaggregation cases, and the orange points stand for the aggregation cases. (a) The $128 \times 128 \text{ km}^2$ domain averaged CWP ($\text{kg} \cdot \text{m}^{-2}$), (b) the maximum of the CWP ($\text{kg} \cdot \text{m}^{-2}$) in the $128 \times 128 \text{ km}^2$ crops. (c) The cloudy area covers $128 \times 128 \text{ km}^2$ crops after being masked by the value of $2 \text{ kg} \cdot \text{m}^{-2}$. 33

3-4 (a) and (b) show the example of the comparison between nonaggregation and aggregation labeled data in the training dataset. The cropped cloud water path field on a $128 \times 128 \text{ km}^2$ is masked by a $2 \text{ kg} \cdot \text{m}^{-2}$ threshold. The green grid will be regarded as one, while dark blue grids will be regarded as zero. The ones and zeros will be inputted into the input layer of the CNN model. The graph in (c) shows the referenced cloud water (masked) inputted into the CNN models. From (d-0) to (d-4), show the graphs of the encoded channel 0 to 4, representing the max-pooling layer with the 42×42 grid or the $128 \times 128 \text{ km}$ domain. Each grid concludes the signal from the $3 \times 3 \text{ km}^2$ box. The brightness in the graphs shows the strength of the signal outputted. Each channel has its scale, so we will only focus on the relative strengths and the spatial distributions. The channel order may vary due to the randomness of the training process; however, the characteristics of the relative strengths and the spatial distributions will remain. 34

4-1 Cold pool and convection characteristics from Large-Eddy simulations. Snapshots from high-resolution Large-Eddy Simulation (LES) illustrating two convective cases, case 1 from (A) to (C), and case 2 from (D) to (F). Outgoing Longwave Radiation (OLR, as shown in A and D) indicates cloud-top temperatures and convective activity with the domain-averaged value on the bottom right. Cold pool intensity (as shown in B and E) computed by vertically integrating negative buoyancy over the boundary 40

layer, and red contours denote precipitation rates exceeding 1 mm h^{-1} . Vertical cross-sections (as shown in C and F) show the vertical structures of cloud water, ice, and rainwater within the boxed regions in (A) (B) and (D) (E). The top row from (A) to (C) represents shallow convection with limited OLR depression and small cold pools, while the bottom row from (D) to (F) shows deep convection with extensive anvil clouds and deep vertical development.

4-2 U-Net framework for cold pool intensity prediction. This figure presents the data processing pipeline and U-Net architecture used to predict high-resolution cold pool intensity from coarse-resolution atmospheric fields. (A) Subsampled Input: Low-resolution input fields (wind, relative humidity, potential temperature) at 0.8, 1.6, 3.2, 6.4 km resolution; we take 3.2 km resolution as an example in this figure. (B) U-Net structure: The deep learning network structures used in this study, which predict high-resolution cold pool intensity from low-resolution physics fields. (C) Cold pool intensity binary masks: The model outputs five binary masks corresponding to cold pool intensity thresholds from 1 to 5 m/s. These are compared with LES ground truth to compute Dice loss and guide model training. (D) Synthesized cold pool intensity and LES ground truth: The binary masks are combined to reconstruct the predicted cold pool field (left), which can be compared visually against the LES ground truth (right). The Dice score shown in the bottom right quantifies the prediction quality. 42

4-3 Comparison of cold pool predictions using U-Net and linear regression models. (A) High-resolution cold pool intensity from LES (100 m resolution). (B) Cold pool boundaries (marked as “Edge” in B and with orange boundary areas in C to J) based on 10-min propagation. (C to F) U-Net predictions at input resolutions of 0.8, 1.6, 3.2, and 6.4 km, and (G) to (J) Linear regression model predictions at the same resolutions. The value D on each subgraph's title represents Dice coefficients, which quantify the prediction accuracy. 45

4-4 Evaluation of U-Net cold pool predictions at 3.2 km resolution. This figure provides a detailed comparison between the cold pool intensity from LES and the U-Net predictions. Figure (A) represents the example of cold pool intensity from LES at 100 m resolution, and (B) the corresponding synthesized U-Net prediction. The corresponding input low resolution 46

physics fields are shown from (C) to (F), we take 3.2 km resolution input fields in this example. From (G) to (K) compare the ground truth cold pool intensity masks (gray regions) with U-Net predicted intensity masks (blue regions) at different intensity thresholds (1.0, 2.0, 3.0, 4.0, and 5.0 m/s). (L) presents a schematic of the intensity mask ranges predicted by LES and U-Net.

4-5 Performance comparison of U-Net and linear regression in predicting cold pools. Boxplots show Dice coefficients for U-Net (green) and linear regression (blue) at four resolutions (R_08: 0.8, R_16: 1.6, R_32: 3.2, R_64: 6.4 km) using subsampled input. At each resolution, the Dice coefficients are computed based on comparisons over 720 test samples. (A) represents the cold pool intensity prediction, while (B) covers the cold pool intensity and boundary prediction. The result shows that U-Net model outperforms linear regression model across all resolutions. 47

4-6 Sensitivity analysis of cold pool prediction using different input fields. Each row shows U-Net input maps (U, V, RH, θ) on 3.2 km resolution (A–D, F–I, and K–N), with the corresponding cold pool intensity from U-Net prediction (E, J, and O) and LES (P). Group 1 from (A) to (E) includes all variables included, group 2 from (F) to (J) masked the dynamic variables (the wind components), and group 3 from (K) to (O) mask the thermodynamic variables (the RH and θ). The ground truth of cold pool intensity from LES is shown in (P). Results show that wind information may be more critical for U-Net cold pool prediction. 47

4-7 Impact of input variables on cold pool prediction performance. Boxplots show correlation coefficients between predicted and ground-truth cold pool fields for four resolutions (R_08 to R_64). Colors represent the combination of input fields: Thermal fields only (pink), dynamic fields only (green), and all variables (blue). The group with the wind fields appears to improve the predictions, especially at finer resolutions. For each resolution and each input group, the correlation coefficients are computed based on 720 test samples. 48

List of Tables



3-1 R values of the cloud fraction/fractal dimension and the channel signal sums corresponded to the model used in Figure 3-4 (d-0) to (d-4), separated by the aggregated and non-aggregated cases. 35

Chapter 1



Introduction

1.1 The Rise of Deep Learning in Atmospheric Sciences

Adapting deep learning techniques to atmospheric science research has become a prominent field in recent years. The use of neural network statistical methods has been widely applied to climatology, meteorology, atmospheric chemistry, and other atmospheric sciences fields (Gentine et al., 2018; Rasp et al., 2018; Yuval & O'Gorman, 2020). By building various neural network frameworks and using both existing observational and simulated data, data-driven prediction model construction can be achieved automatically.

As neural network architectures become increasingly complex, deep learning, which involves increasing the number of parameters or deep architectural depth in neural networks, has been introduced to atmospheric sciences research. In the past few years, with the development of artificial intelligence receiving significant attention and resources in the data science field, the application of deep learning techniques—especially the generative model—in atmospheric forecasting has increased. These applications range from down-scaling and probabilistic forecasting to satellite image recognition (Agrawal et al., 2019; Brüning et al., 2024; Higa et al., 2021; Hoeller et al., 2022).



1.2 The Problem of Black Box: The Need for Interpretability and Explainability

With high-quality datasets, these deep learning models often achieve acceptable results in prediction tasks. However, their output relies on computations among numerous parameters, making it difficult for researchers to understand why specific results are produced. This deep learning model operates like a "black box," which may lack transparency, interpretability, and explainability (Samek et al., 2017).

These black box models often cannot explain the causal relationships between predictions and the physics variables. The increasing number of high-accuracy black box models raises significant challenges for scientific discussions, and there are many doubts about the robustness and generalizability of these models (Beucler et al., 2021). Without a clear understanding of the underlying reasoning, applying these models to study complex atmospheric dynamics remains risky and scientifically insufficient.

1.3 Morphology as a Physical Indicator: Linking Visual Features to Domain Knowledge

To address this lack of underlying reasoning, we propose a framework focused on specific physical attributes. In this research, we tried to construct the linkage between the morphology of atmospheric physics variables and the deep learning model predictions. We chose morphology as an important indicator because it not only integrates the outcomes of nonlinear physical processes, but also provides interpretable visual cues. Human visual perception plays an important role in judging the reasonableness of natural

variables, including their correlations and causal inferences.



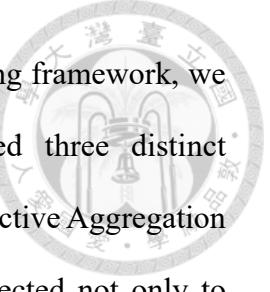
When interpreting deep learning models, we can apply our domain-knowledge of atmospheric physical variables to effectively assess the validity of the outputs and reasoning processes of the deep learning models. Using the morphology of physics variables to assess the reasonableness of deep learning model construction can be applied to various scales of atmospheric science problems.

In this research, we hypothesize that by extracting nonlinear morphological features, we can bridge the gap between data-driven predictions and physical consistency. This approach is tested across three distinct scales of atmospheric processes:

1. **Large-scale:** The changes in the morphology of pressure fields during sudden stratospheric warming (SSW) phenomena.
2. **Mesoscale:** The relationship between aggregated convection development characteristics and cloud water distribution morphology.
3. **Small-scale:** The relationship between the cold pool in the subgrid-scale of Global Storm Resolving Model (GSRM) and the environmental variables morphology.

1.4 Thesis Objectives and Framework

Convolutional neural networks (CNNs) and their variants are frequently used to capture the morphology of physical variables and construct predictive models (Chattopadhyay et al., 2020; Ham et al., 2019; Weyn et al., 2020). By using CNNs, we can allow the model to automatically determine and extract key features such as edges, areas, positions, maxima, and minima, and even complex nonlinear structures from the input data.



To demonstrate the robustness and generalizability of this deep learning framework, we adopted a cross-scale approach for case selection. We identified three distinct phenomena—Sudden Stratospheric Warmings (Planetary scale), Convective Aggregation (Mesoscale), and Cold Pools (Sub-grid scale). These cases were selected not only to represent a hierarchy of atmospheric scales but also because they share a common scientific challenge: their physical states and evolutions are governed by nonlinear spatial structures (or morphological features) that are difficult to quantify using traditional scalar statistics. By applying our framework across these diverse scales, we aim to validate that morphology can serve as a universal physical indicator connecting data patterns to dynamic mechanisms.

However, the specific morphological tasks differ across these scales, ranging from global evolution representation to local feature extraction and subgrid reconstruction. Therefore, various CNN-based network architectures must be adapted. Our research used three different CNN-based deep learning networks to build the framework linking physical variables' morphology in atmospheric science problems of various scales. This dissertation contains the following three parts:

1. Representing Large-Scale Evolution (Chapter 2):

In the first part of this study, we conducted a morphological analysis of sudden stratospheric warming (SSW) phenomena. Traditional classification methods (e.g., linear PCA) may not be suitable to categorize all events fully due to the complex spatial deformation of the polar vortex. Therefore, we demonstrated the use of a combination of Variational Autoencoder (VAE) and CNN to capture the large-scale atmospheric physics variables and compress the morphology into scatter plots. This visualization

effectively captures the temporal and spatial changes of the 10 hPa geopotential height, providing a continuous phase diagram to study the spatial-temporal characteristics of SSWs.



2. Extracting Mesoscale Nonlinear Features (Chapter 3):

In the second part, we investigate the relationship between cloud water distribution patterns and the development of aggregated convection. We constructed a framework that uses CNN models and gradually reduces the key physics features of the input dataset. By repeatedly removing characteristics we believed were important (such as cloud water path magnitude) and retraining the model, we discovered the nonlinear relationships between edge complexity and the area of cloud water patterns. This "iterative feature removal" strategy allows us to identify the important morphological features that distinguish convective aggregation.

3. Reconstructing Subgrid-Scale Dynamics (Chapter 4):

In the third part, we explore the potential of deep learning as a subgrid parameterization tool. We use a CNN-based U-Net deep learning model to predict high-resolution cold pool patterns from low-resolution physics fields—mimicking the resolution gap in Global Storm-Resolving Models (GSRMs). By designing the U-Net output as binary masks of cold pool intensity, we can focus on the overall range and intensity distribution pattern. Analyzing the key morphology concluded by U-Net enables us to understand the important physics variables (specifically dynamic fields vs. thermodynamic fields) required for parameterizing high-resolution cold pools.

These chapters are derived from three scientific papers conducted during the doctoral

program, each corresponding to a specific peer-reviewed publication:

Chapter 2 is adapted from the work published in Chen et al. (2024), "Exploiting a variational auto-encoder to represent the evolution of sudden stratospheric warmings."

Environmental Research: Climate 3(2), <https://doi.org/10.1088/2752-5295/ad3a0d>, which investigates the morphological evolution of Sudden Stratospheric Warmings using a variational autoencoder framework.

Chapter 3 corresponds to the study presented in Chen et al. (2023), A deep learning framework for analyzing cloud characteristics of aggregated convection using cloud-resolving model simulations. Atmospheric Science Letters, 24(5), e1150. <https://doi.org/10.1002/asl.1150>, focusing on extracting nonlinear features of mesoscale convective aggregation via an iterative deep learning approach.

Chapter 4 is based on Y. C. Chen and C. M. Wu (2025). "Capturing Subgrid Cold Pool Dynamics With U-Net: Insights From Large-Eddy Simulation for Storm-Resolving Modeling." Atmospheric Science Letters 26, no. 7: e1309. <https://doi.org/10.1002/asl.1309>, demonstrating the reconstruction of subgrid-scale cold pool morphology from coarse-grained environmental fields.

Following these case studies, Chapter 5 synthesizes the findings across these three distinct scales, providing a general discussion on the role of different neural network architectures and the interpretability of morphology as a physical indicator. Finally, Chapter 6 summarizes the major scientific contributions of this dissertation and outlines potential directions for future research.

Through these parts of studies, this dissertation aims to establish an explainable deep learning framework that not only achieves acceptable prediction accuracy but also reveals the underlying nonlinear morphological features of atmospheric processes.



Chapter 2



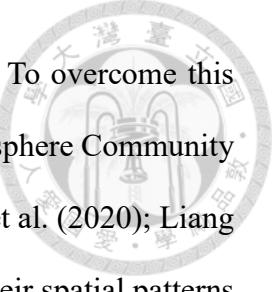
Representing the Morphological Evolution of Large-Scale Circulation using Variational Autoencoder

2.1 Introduction

Sudden stratospheric warmings (SSWs) are the most dramatic events in the wintertime stratosphere, characterized by substantial disruption to the stratospheric polar vortex (Andrews, 1987; Baldwin et al., 2021; Charlton & Polvani, 2007; Limpasuvan et al., 2004). These extreme events are typically categorized into displacement and splitting types depending on the morphology of the vortex (Baldwin et al., 2021; Charlton & Polvani, 2007; Lehtonen & Karpechko, 2016; Matthewman et al., 2009; Mitchell et al., 2011; Seviour et al., 2013). Monitoring the genesis and evolution of these morphological changes is important because they are usually followed by anomalous tropospheric circulation regimes that are important for subseasonal-to-seasonal prediction (Baldwin et al., 2003; Davis et al., 2022; Domeisen et al., 2020; Kidston et al., 2015; Sigmond et al., 2013; Tripathi, Baldwin, et al., 2015; Tripathi, Charlton-Perez, et al., 2015).

Despite several analysis methods that have been used to study the evolution of SSWs, the potential of deep learning methods to capture these complex spatial structures has not yet been fully explored. This is mainly due to the relative scarcity of observed events, which

creates a small sample size problem for training data-driven models. To overcome this limitation, we use data from historical simulations of the Whole Atmosphere Community Climate Model version 6 (WACCM6, Gettelman et al. (2019); Liang et al. (2020); Liang et al. (2022)) to identify thousands of simulated SSWs. We then use their spatial patterns to train a deep learning model.



In this chapter, we use a convolutional neural network (CNN) combined with a variational auto-encoder (VAE, Kingma and Welling (2013))—a generative deep learning model—to construct a phase diagram that characterizes the SSW evolution. This approach aligns with our dissertation's goal of extracting nonlinear features: it allows us to create a latent space that encapsulates the essential features of the vortex structure during SSWs. By projecting the high-dimensional geopotential height fields into a compressed 2D space, this visualization effectively concludes the temporal and spatial changes of the polar vortex. Unlike linear methods such as Principal Component Analysis (PCA), the VAE captures the nonlinear continuous transition of the vortex pattern, providing a new perspective for discussing the evolutionary paths of SSWs.

2.2 Data and Methods

2.2.1 Data Description and Preprocessing

We use daily-mean zonal wind and geopotential height fields at 10 hPa (Z10) from the National Aeronautics and Space Administration's Modern-Era Retrospective Analysis for Research and Applications, version 2 (MERRA-2, Gelaro et al. (2017)) reanalysis and two sets of large-ensemble historical simulations conducted using WACCM6. The WACCM6 simulations provide a total of 1539 SSW events, offering a rich dataset to train

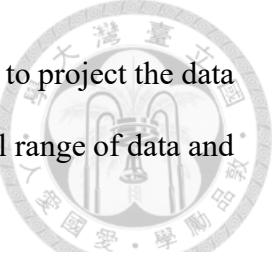
our deep learning model.



The WACCM6 is a high-top chemistry-climate model with 70 vertical levels extending from the surface to $6 \times 10^{-6} \text{ hPa}$ ($\sim 140 \text{ km}$), allowing for a robust representation of stratospheric dynamics. The simulations have a horizontal resolution of $0.95^\circ \times 1.25^\circ$ in latitude and longitude. The dataset comprises two sets of historical simulations (1979–2014) forced by observed sea surface temperatures and sea ice concentrations, with each set containing 30 ensemble members generated via micro-perturbations to the initial atmospheric temperature fields. This large-ensemble approach (totaling 2,160 simulation years) is important for overcoming the sampling limitations of observational records when studying extreme events like SSWs.

Specifically, the simulations consist of two configurations to account for different forcing factors. The first set involves WACCM6 being forced by the daily time-varying global sea-surface temperature (SST) and sea-ice concentration (SIC) (Eyring et al., 2016; Haarsma et al., 2016). The second set replicates the first but replaces the Northern Hemisphere SIC with daily climatological values. This experimental design ensures that the atmospheric circulation in the second set remains unaffected by variations in Arctic sea ice, allowing us to isolate the stratospheric variability driven by internal atmospheric dynamics from that forced by sea-ice anomalies.

As discussed in Chapter 1, appropriate preprocessing must be done according to the data characteristics to facilitate the CNN's ability to extract key features. The Z10 field output from WACCM6 simulations is originally in a latitude-longitude grid. To maintain the grid area size at high latitudes and avoid distortion that might confuse the morphological



learning, we employ a bilinear interpolation method (Kim et al., 2019) to project the data onto a generic square grid. Figure 2-1 shows an example of the original range of data and the interpolated data.

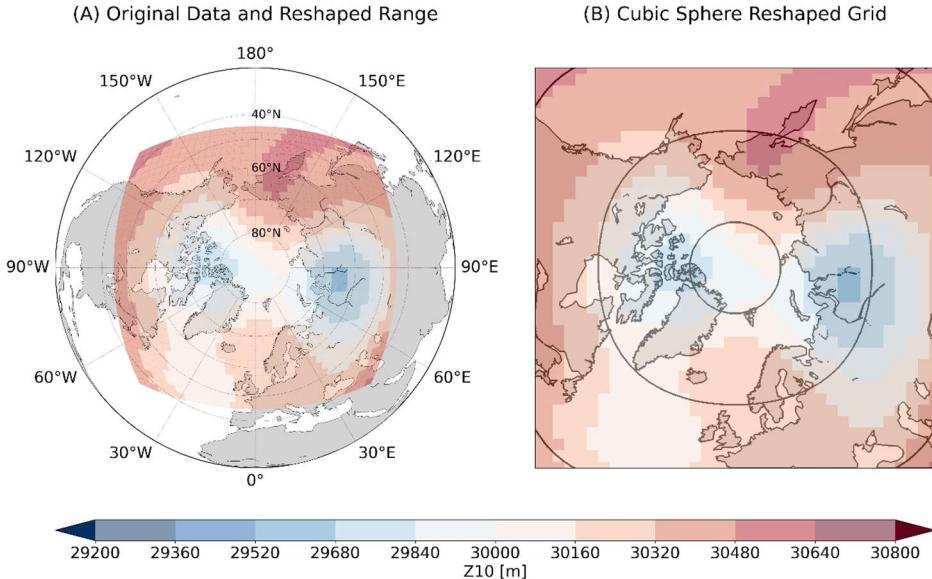


Figure 2-1. Illustration of the bilinear interpolation method applied to a Z10 field of a split-type SSW event from WACCM6 simulations. (a) The longitude-latitude map of the Northern Hemisphere in an orthographic projection with 90° N at the center, with the region of color shadings indicating the domain to be interpolated. (b) The resultant 36 × 36 squared grid after using the bilinear interpolation method. The parallels on are respectively at 80, 60, and 40 degrees from the innermost to the outermost.

This spatial interpolation allows the VAE model to extract features in a more feasible way (i.e., squared images) as the stratospheric polar vortex can be distorted in the longitude-latitude grid. These squared images of Z10 fields are then standardized to the Z-score to accelerate the convergence rate of the neural network.

2.2.2 Definition and Classification of SSWs

To identify SSW events, we follow the standard criterion proposed by Charlton and Polvani (2007), defined by the reversal of zonal-mean zonal winds at 10 hPa and 60°N from westerly to easterly during the boreal extended winter (November to March). The onset date (day 0) is defined as the first day of wind reversal.

Furthermore, to categorize the morphology of the polar vortex, we adopt the two-dimensional vortex moment analysis (Seviour et al., 2013). This method calculates two key geometric indices based on the geopotential height field: the aspect ratio (r) and the centroid latitude (ϕ_c).

Displacement Events: Defined when the vortex centroid is shifted equatorward ($\phi_c < 66^\circ N$) for at least 7 days, while the aspect ratio remains low.

Splitting Events: Defined when the vortex becomes elongated and splits, indicated by an aspect ratio ($r > 2.4$) for at least 7 days.

Using these criteria, we identified 1,177 displacement events and 362 splitting events from the WACCM6 simulations.

2.2.3 Variational Auto-Encoder (VAE) Framework

To capture the morphology of the polar vortex, we constructed a CNN-based VAE model. The VAE consists of two components: an encoder and a decoder. The encoder allows the model to automatically determine and extract key spatial structures—such as the vortex edge, position, and intensity—and compresses this information into a latent space. The decoder then samples from this compacted representation to generate (reconstruct) the output data. The structure of CNN-based VAE model used in this study can be found in Figure 2-2.

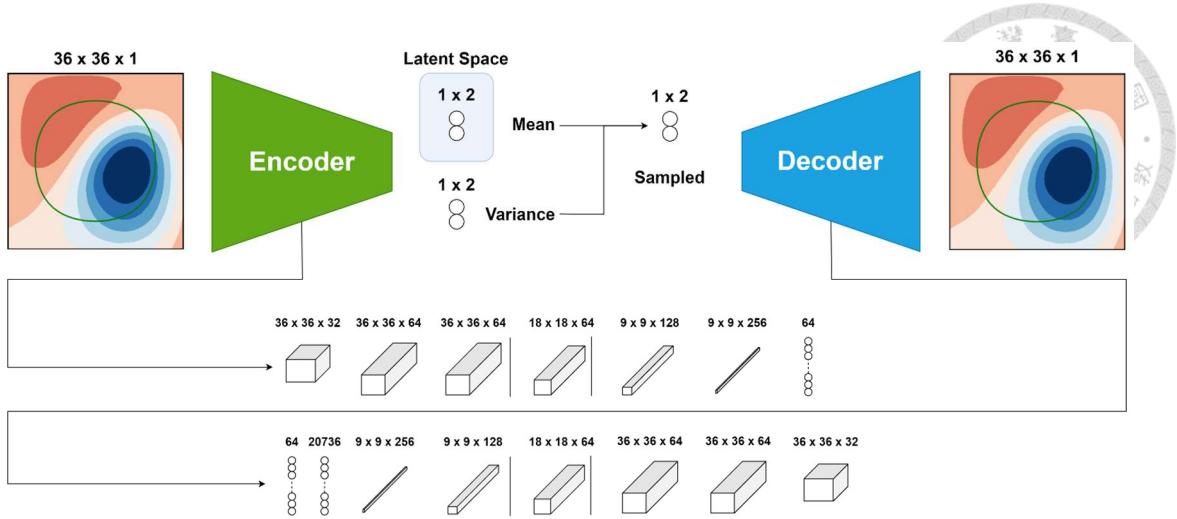


Figure 2-2. A schematic of the VAE structure used in this study. The top row represents the combination of a decoder and an encoder, which form the VAE model. The middle row illustrates the details of the encoder, while the bottom row those of the decoder. The vertical solid line after each layer denotes the max-pooling/upsampling layers, which reduce or increase the dimension of pixelated images, whereas the cubes represent the convolutional layers with varying numbers of channels. The encoder reads in a 36×36 Z10 input image, compresses its information through the convolutional layers, and outputs two means (μ_0, μ_1) and two variances (σ_0^2, σ_1^2) that form the latent space. The decoder utilizes these sampled results based on Gaussian distribution and generates a 36×36 Z10 output through multiple convolutional layers. The mean squared error (MSE) between the original and generated images and the Kullback–Leibler divergence between the latent space and the Gaussian distribution are computed. These errors are backpropagated to train the encoder and decoder. The light blue color shading highlights the two means of the latent space, which is used to produce the phase diagram in this study.

Looking into the latent space yields an understanding of how the VAE learns the continuous variations in the patterns. In this study, we focus on the two "modes" (dimensions) of the trained VAE, z_0 and z_1 . These two dimensions form a two-dimensional phase diagram, making it easy to examine the evolution of the SSWs. By minimizing the difference between the input and the generated output, the VAE optimizes the latent space to represent the continuous variations of the vortex patterns from the input data.

To provide a baseline for assessing the VAE's ability to capture nonlinear morphology, we also perform Principal Component Analysis (PCA) over the same Z10 training data.

As a linear baseline, the first two principal components (PCs) explain 30.6% and 21.4% of the total variance, respectively, accounting for over 50% of the variability combined.

While PC1 typically represents the vortex displacement and PC2 captures the wave-number 1 or 2 features, we will demonstrate in the following sections that this linear decomposition struggles to fully capture the complex, non-linear deformation of splitting events compared to the VAE. Physically, the first PCA mode is characterized by a dipole in Z10 anomalies between northwestern North America and the Barents-Kara Seas, while the second mode features lower Z10 values over Greenland together with higher values over eastern Siberia. Although these two modes are statistically distinguishable based on North's rule (North et al., 1982) and combined explain more than 50% of the total variance, their linear combination struggles to represent the continuous structural deformation observed during splitting events.

Before constructing the balanced dataset, we processed the entire collection of identified SSW events. We prepared a total of 32,319 Z10 images from the Z10 anomalies over 10 days prior to and after the onset date. To ensure robust model training and evaluation, we adopted an 80%–9%–11% ratio for separating the data, resulting in 25,893 images for training, 3,024 for validation, and 3,402 for testing. This strict separation ensures that the morphological features learned by the VAE are generalizable and does not result from overfitting to specific events.

To prevent the model from being biased toward the more frequent displacement events (class imbalance), we constructed a balanced training dataset. We randomly sampled an equal number of images (1,299 images) from three categories: displacement, splitting, and 'neither' (transition states), resulting in a total of 3,897 images for training. The model

was trained using the Adam optimizer.



2.3 Results: The Morphological Phase Diagram

A key outcome of the VAE is the generative phase diagram obtained by passing varying latent variables (z_0, z_1) through the decoder. This diagram provides insight into the morphology and spatiotemporal evolution of SSWs. As shown in Figure 2-3, we can observe that the phase diagram captures a wide variety of various vortex structures.

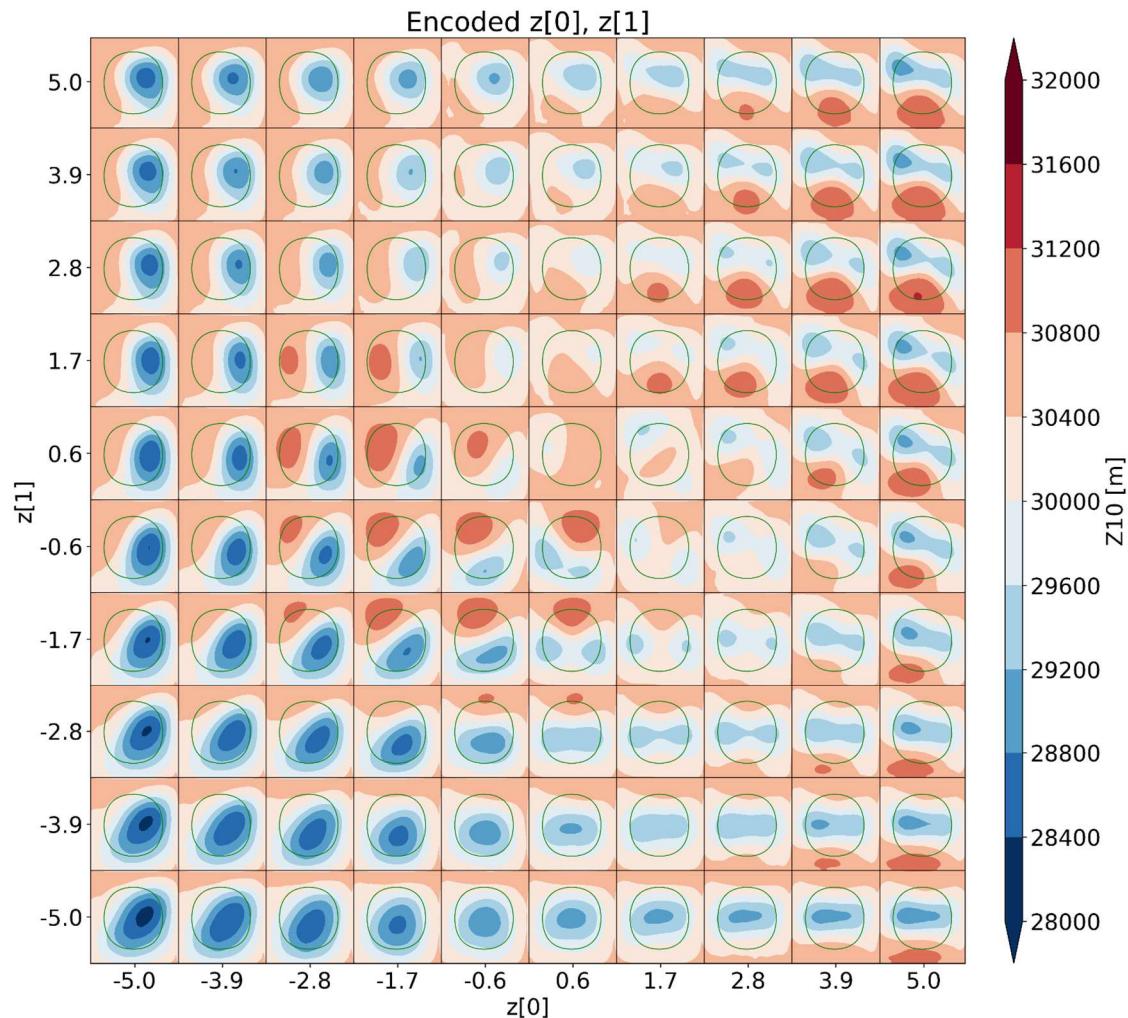
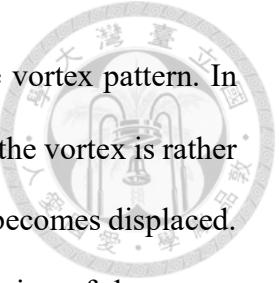


Figure 2-3. This figure illustrates the corresponding $Z10$ field generated by the VAE decoder when given the varying values of z_0 and z_1 in the latent space. The green circle in each panel represents the circle of $60^\circ N$.



The constructed phase diagram depicts a continuous transition of the vortex pattern. In the upper-left quadrant (large negative z_0 , positive z_1) in Figure 2-3, the vortex is rather circular and centered near the pole. Towards the lower-left, the vortex becomes displaced. In the upper-right quadrant, the spatial pattern demonstrates the splitting of the vortex into two smaller vortices.

Interestingly, the phase diagram shows a gradual, continuous transition of the vortex structure from one shape to another. This confirms that the VAE has learned to map the complex, nonlinear morphology of the stratospheric vortex onto a simplified 2D vector.

2.3.1 Evolutionary Tracks and Reconstruction

To further shed insight on the capability of the VAE phase diagram, we project the temporal evolution of specific SSW events onto the diagram. Visualization in Figure 2-4 allows us to track the morphological changes through the 2D phase-space.

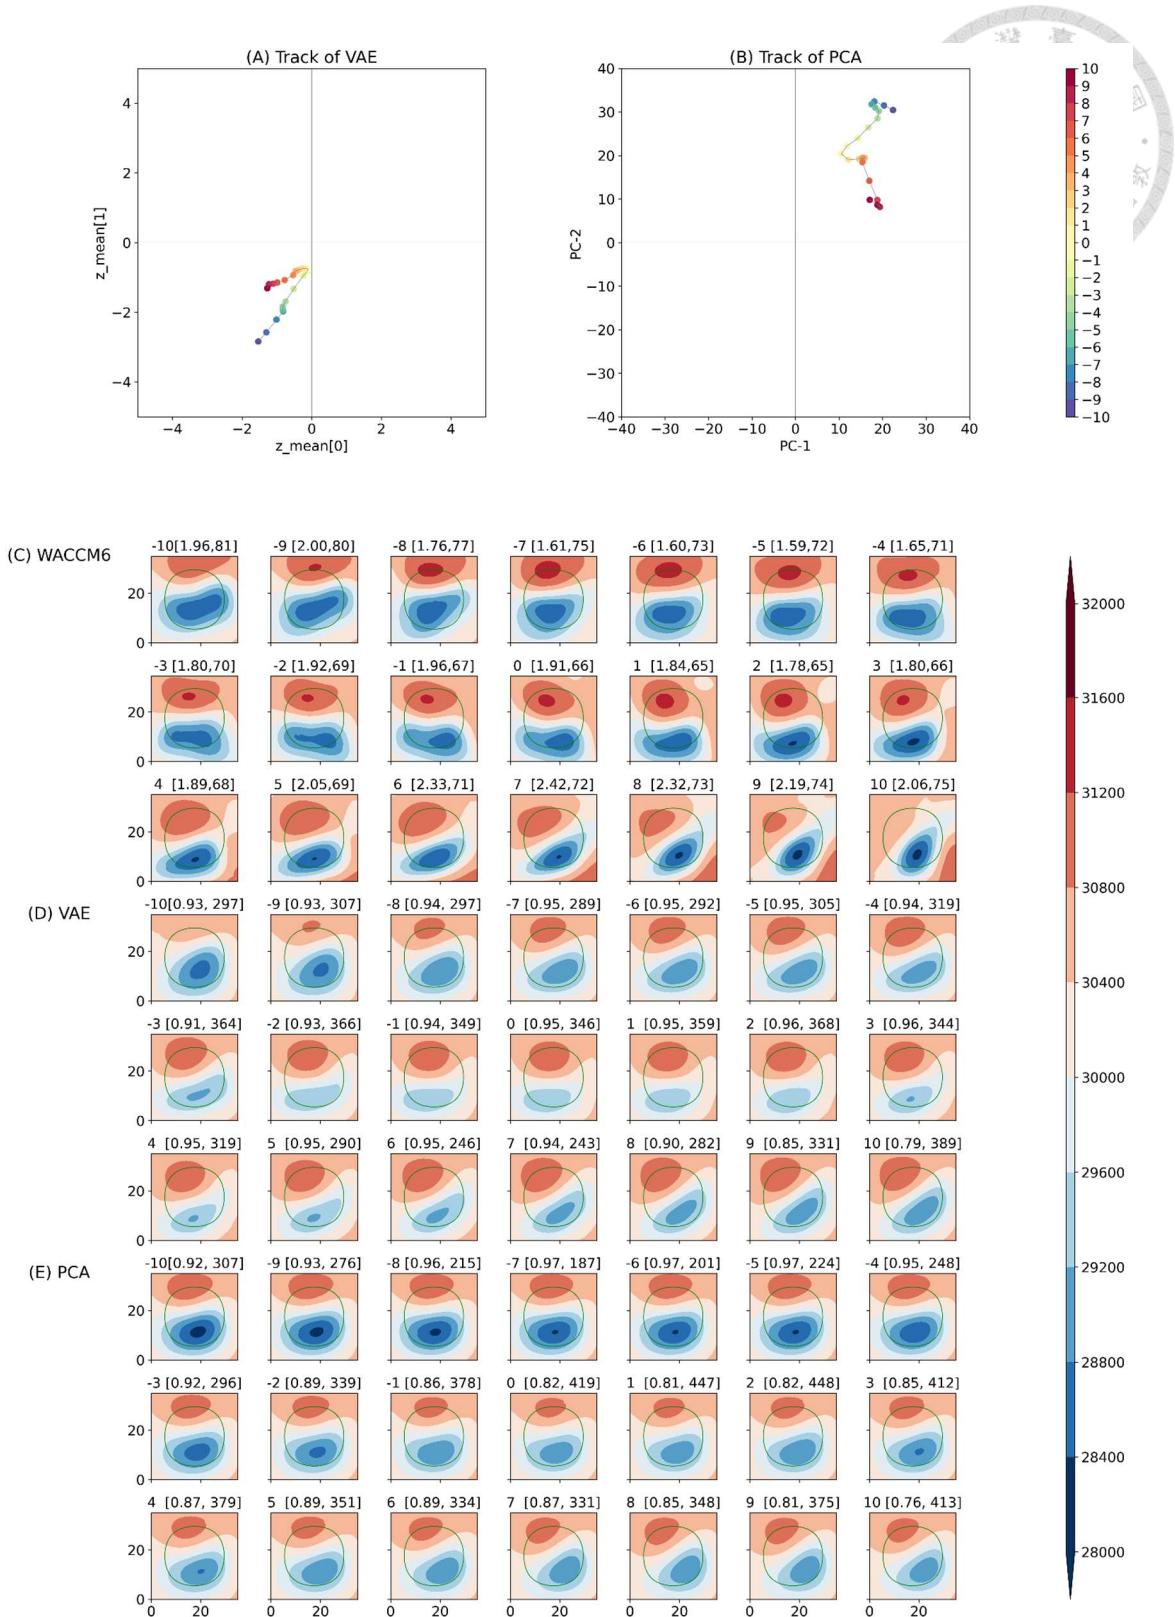


Figure 2-4. (a) The evolution track of a displacement-type SSW event from WACCM6 simulations mapped onto the VAE phase diagram. (b) As in (a), but for the track mapped onto the PCA phase diagram. (c) The WACCM6 Z10 spatial patterns during days [-10, 10]. (d) The VAE reconstructed Z10 fields during days [-10, 10]. (e) As in (d), but for PCA reconstructed Z10 fields. The values in the subtitle parenthesis in (c) are the aspect ratio (r) and central latitude (ϕ_0) of this SSW event, whereas those in (c) and (d) denote the

pattern correlation (R) and RMSE. The green circle in each panel of (c)–(e) represents the circle of $60^\circ N$.

Figure 2-4 shows a displacement-type event. The VAE track stays within the lower-left quadrant, consistent with the displacement-dominant status of the event. The VAE reconstructed fields (Figure 2-4D) successfully capture the elongated and displaced nature of the vortex. In contrast, while PCA (Figure 2-4E) also captures the general evolution, the VAE gives a better-reconstructed vortex morphology, particularly in preserving the structural integrity of the vortex core.

This advantage of VAE becomes even more apparent for split-type events in Figure 2-5.

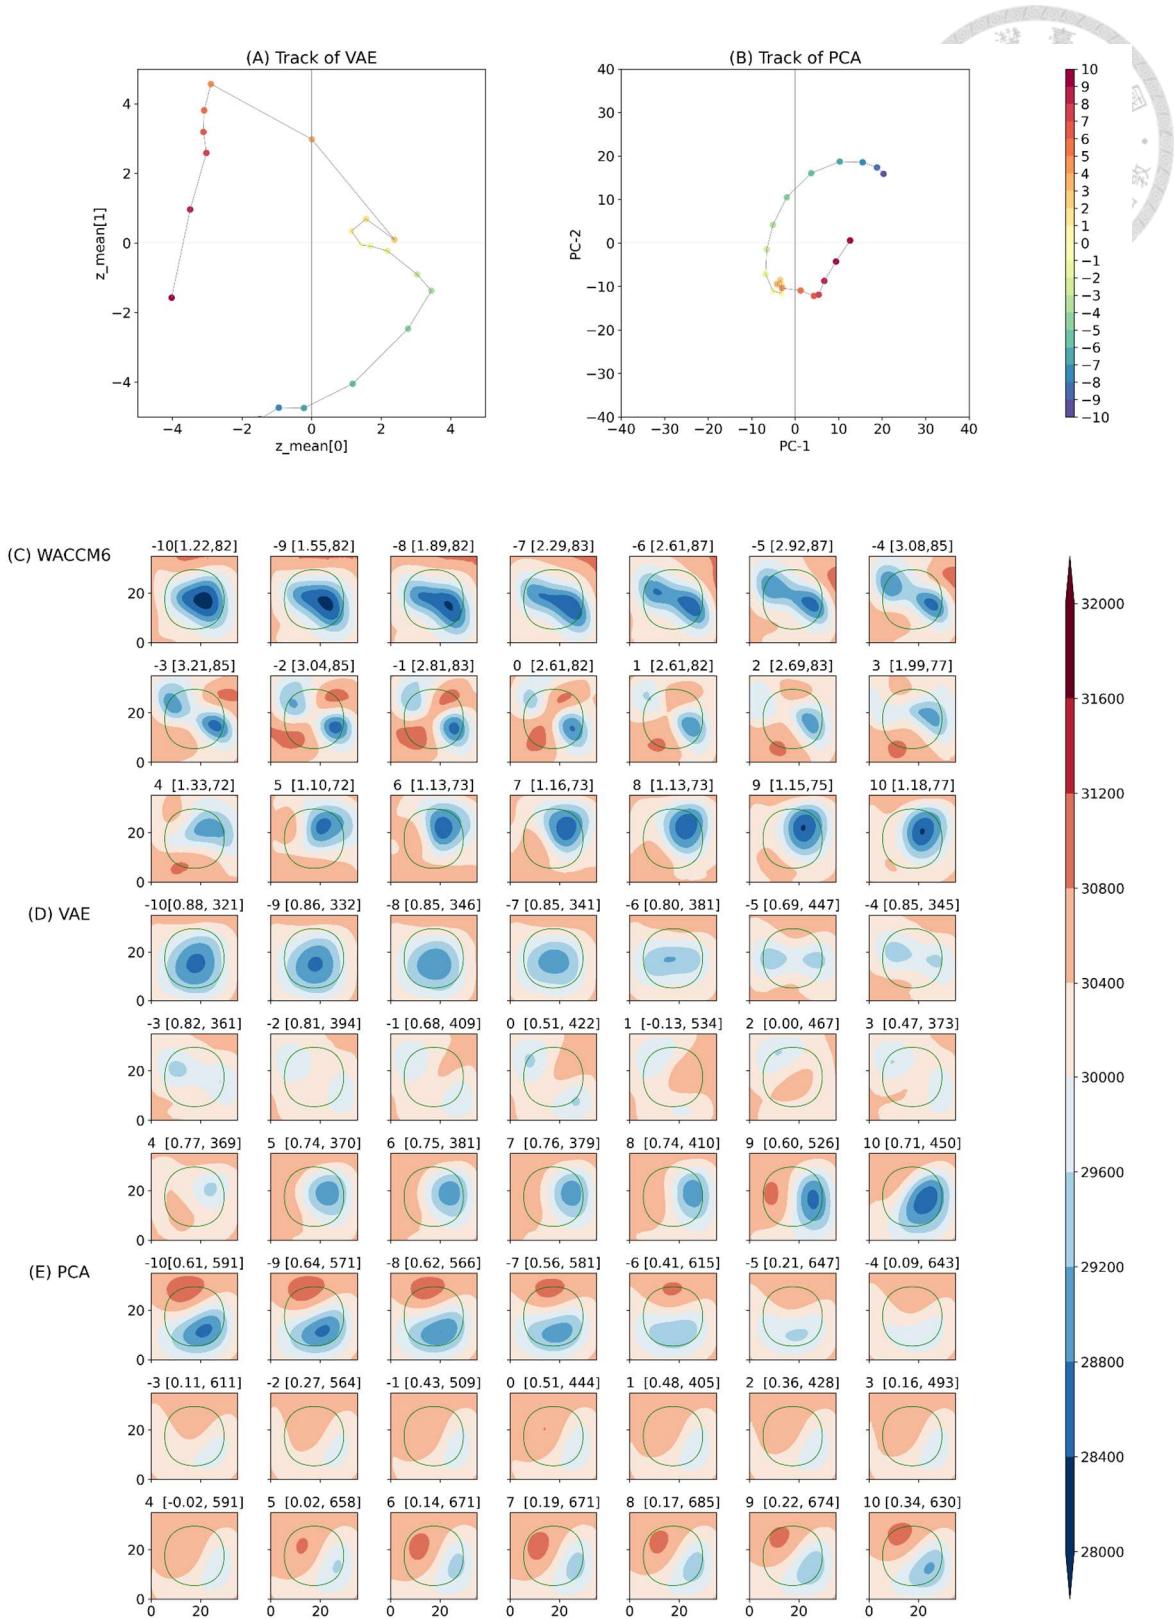


Figure 2-5. The same as in Figure 2-4, but for a split-type SSW.

Figure 2-5 shows a splitting event. VAE is capable of reconstructing the varied spatial patterns—from a circular structure to a stretched one and finally to splitting patterns.

However, the PCA (Figure 2-5E) fails to retrieve these important vortex features; it cannot produce two distinct vortices but instead yields only one vortex. This implies that the linear combination of two PCA modes is insufficient to represent the highly nonlinear morphology of a splitting vortex.

2.3.2 Statistical Distribution and Model Validation

We further investigate if different types of SSWs favor certain regions of the 2D phase-space by mapping all test events onto the heatmap (Figure 2-6).

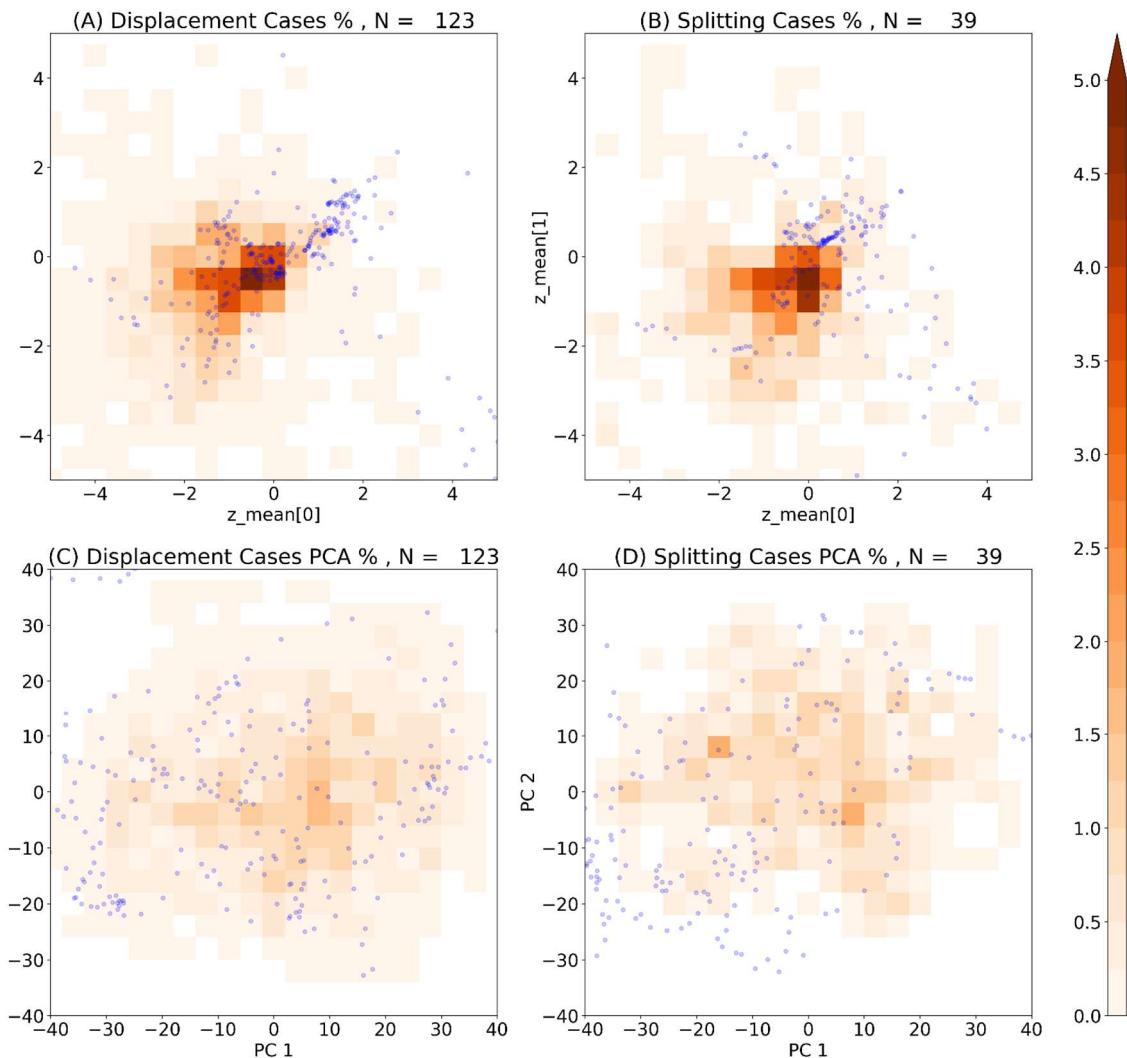


Figure 2-6. (a) The occurrence probability of the Z10 fields associated with all displacement-type SSWs from the WACCM6 testing dataset in the VAE phase diagram. The values are presented as percentages after normalized by the total number of events. (b) As in (a) but for split-type SSWs from WACCM6

simulations. Figures (c) and (d) are the same as (a) and (b), but for the PCA phase diagram. The blue dots are the Z10 fields from the reanalysis dataset mapped onto the VAE and PCA phase diagrams.

The VAE heatmaps (Figures 2-6A, 2-6B) separates the regimes for displacement (lower-left) and splitting (upper-right) events. The observed SSWs from MERRA-2 (blue dots) also map differently onto these regimes, suggesting that the VAE trained on simulation data can be applied to observational data. Conversely, the PCA phase diagram (Figures 2-6C, 2-6D) shows less organized regimes, confirming that VAE offers a more robust framework for mapping vortex structures into two modes.

2.3.3 Quantitative Assessment

To provide quantitative evidence that the VAE reconstructs better evolutions of SSWs than the PCA, we calculate the pattern correlation (R) and root mean squared error ($RMSE$) as a function of time.

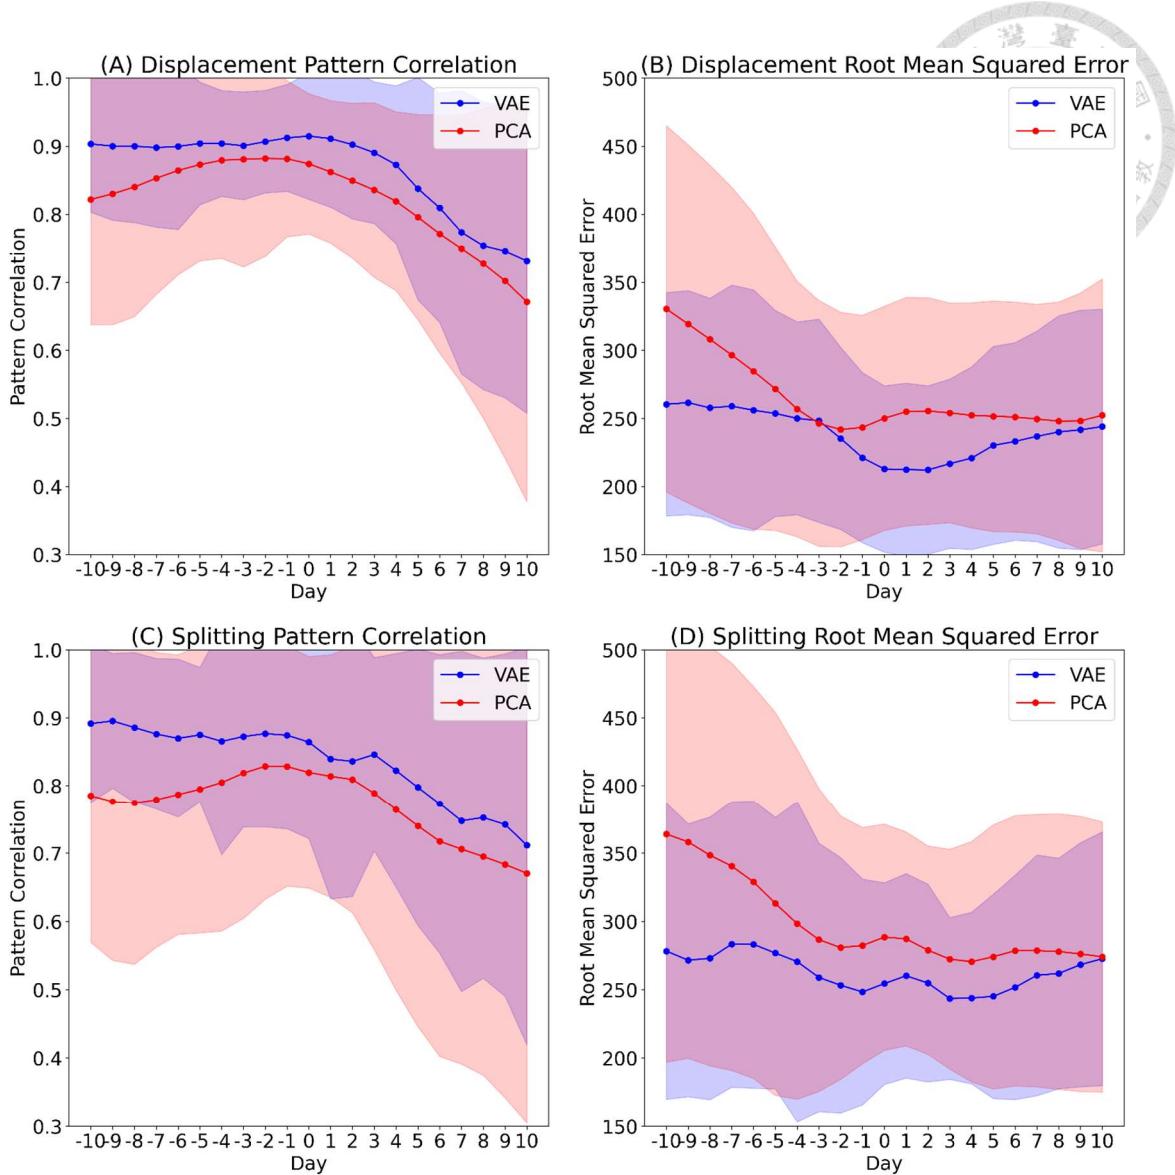


Figure 2-7. The pattern correlation (R) and root mean squared error (RMSE) for (a) and (b) the reconstructed displacement-type SSWs and (c) and (d) the reconstructed split-type SSWs from WACCM6 simulation against corresponding full Z10 during the period of days $[-10, 10]$. The color shading indicates the range of one standard deviation.

Figure 2-7 shows that for both displacement and splitting events, the VAE gives higher R and lower $RMSE$ than the PCA throughout most of the period. This quantitatively supports our visual finding that the VAE performs better in terms of spatial reconstruction.

Finally, we examine the "travel distance" in the phase space to understand the variability

of the events.

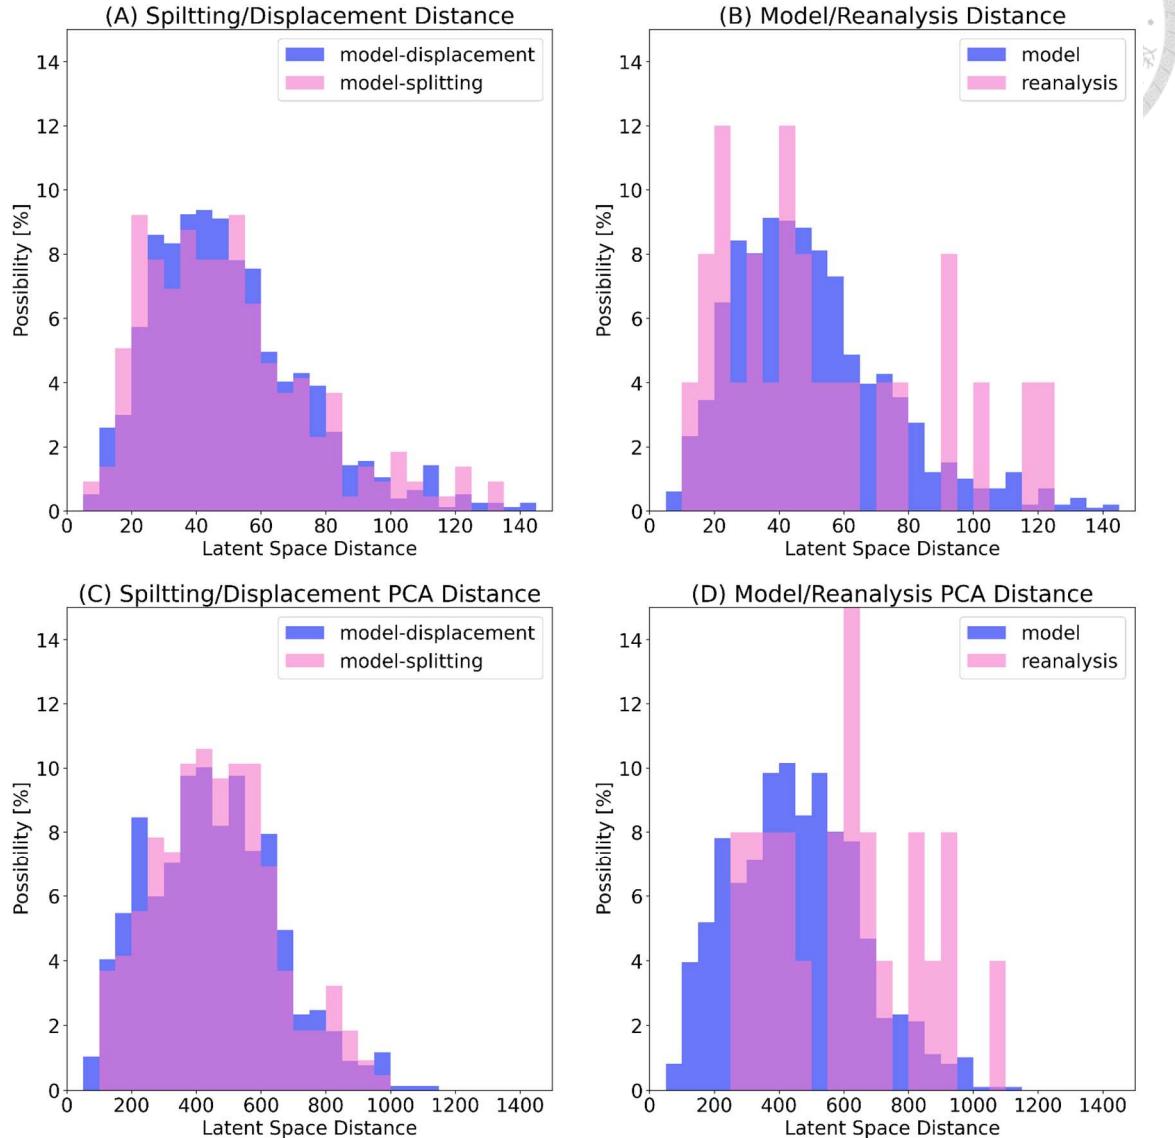


Figure 2-8. (a) Distributions of distance traveled in the VAE phase diagram for all split-type (magenta bars) and displacement-type (blue bars) SSWs. (b) Distributions of distance traveled for all simulated SSWs (blue bars) and observational (magenta bars) SSWs. (c) and (d) The same as (a) and (b) but for the distributions of distance traveled in the PCA phase diagram.

Figure 2-8 shows that the VAE produces a more skewed distribution of travel distances compared to PCA, reflecting its ability to generate more disrupted (and thus more variable) vortex structures.

2.4 Discussion

The latent space of the VAE plays an essential role in this study. We have shown that the dimensions of the latent space (z_0, z_1) are not just abstract numbers but correspond to physical morphological characteristics: z_0 largely controls the transition between displacement (wavenumber-1) and splitting (wavenumber-2) structures, while z_1 relates to the vortex strength and distortion.

This interpretability is practical. It moves the deep learning model away from being a "black box" towards a "glass box" where we can observe the reasoning process. The phase diagram provides an innovative means to study the evolution of the vortex structure in the duration of SSW, revealing that splitting events are not just discrete jumps but follow continuous morphological pathways.

Furthermore, the VAE phase diagram provides an innovative approach to manifest the discrepancy between SSWs in global climate model simulations and reanalysis data. As shown in the distribution comparison (Figure 2-6), the simulated vortex structures from WACCM6 occur less frequently in the upper-right quadrant compared to the observed events from MERRA-2. This implies that there may exist model biases in the stratospheric polar vortex of WACCM6, particularly affecting the structure of stationary waves. By projecting both datasets onto the same latent space, the VAE effectively highlights these subtle morphological differences that scalar metrics might miss, demonstrating its potential as a benchmark tool to evaluate model performance in simulating extreme stratospheric events.

However, a question arises: do the generated vortex structures in the latent space (e.g., in the empty quadrants of the phase diagram) represent physically realizable states? While

some patterns may be rare in the current climate, they represent potential morphological configurations that the atmosphere could theoretically explore. This generative capability is a distinct advantage of the VAE over discriminative models.



2.5 Summary and Conclusion

In this chapter, we have successfully trained a VAE using the stratospheric polar vortex during SSWs from large-ensemble WACCM6 simulations. By deriving a phase diagram from the latent space, we characterized the spatiotemporal evolution of SSWs from a morphological perspective.

Compared with the baseline results obtained from the PCA approach, reconstructions from VAE have better spatial structures, showing the vortex stretching, deformation, and splitting with higher fidelity. Also, VAE provides a continuous latent space that maps the nonlinear transition of the vortex.

These results demonstrate that identifying morphological features via deep learning can provide new insights into large-scale atmospheric dynamics. The VAE effectively compressed the global morphology of the polar vortex into a low-dimensional space, allowing us to track its evolution.

Though VAE successfully captured the evolution of the large-scale field by compressing the entire image into a latent vector by encoding the entire map at once, in many atmospheric problems, particularly at the mesoscale, we are interested in identifying which specific local features within the scene drive a physical process (such as convection aggregation). The VAE's latent variables (z_0, z_1) describe the "what" (the

state), but they do not explicitly isolate the "why" (the causal features) in a way that allows us to analyze the physics of organization.



Therefore, in the next chapter, we will shift our focus from representation to feature extraction and explanation. We will introduce an iterative framework to identify the nonlinear morphological features, moving deeper into the interpretability of the neural network's hidden layers.

Chapter 3



Extracting Nonlinear Morphological Features of Mesoscale Convection via an Iterative Deep Learning Framework

3.1 Introduction

Tropical deep convection can aggregate into organized systems that can reach mesoscale or even hundred-kilometer scales horizontally. These aggregated convection systems play an important role in the climate system and regional extreme weather (Hamada et al., 2014). While the theoretical understanding of aggregation has advanced (e.g., the "moist-convective quasi-equilibrium" hypothesis, Arakawa (2004); Arakawa and Schubert (1974)), defining the state of aggregation often relies on simple linear metrics or scalar statistics. However, aggregated convection in nature involves high-dimensional nonlinear relationships of multiple physical characteristics (Xu et al., 2019) that are best observed through their morphology.

As discussed in Chapter 1, identifying these morphological features is challenging for traditional statistical methods but is a strength of Convolutional Neural Networks (CNNs). CNNs can statistically capture the characteristics of gridded data (Chattopadhyay et al., 2020; Higa et al., 2021; Tsou et al., 2019), effectively learning from the "snapshot" of aggregated convection in the model simulations. However, a standard CNN operates as a

"black box." Though it may achieve high accuracy, the specific morphological features it uses to distinguish aggregated from non-aggregated convection remain hidden in the high-dimensional weighted parameters (Ribeiro et al., 2016; Roscher et al., 2020).

In this chapter, we construct an "explainable framework" to extract these hidden relationships. Unlike Chapter 2, where we used a VAE to represent the evolution of a state, here we use a CNN to analyze the state. By iteratively training the model and eliminating known physical features (a process we term "Iterative Feature Removal"), we force the neural network to reveal the nonlinear morphological features that are important for identifying organized convection.

3.2 Methodology

3.2.1 CNN Model and the Hidden Layer

The convolutional neural network (CNN) is a robust model for recognizing the image or the distribution with spatial characteristics. The structure of the CNN applied in this study is visualized in Figure 3-1, which is inspired by the Visual Geometric Group (VGG)-net (Simonyan & Zisserman, 2015).

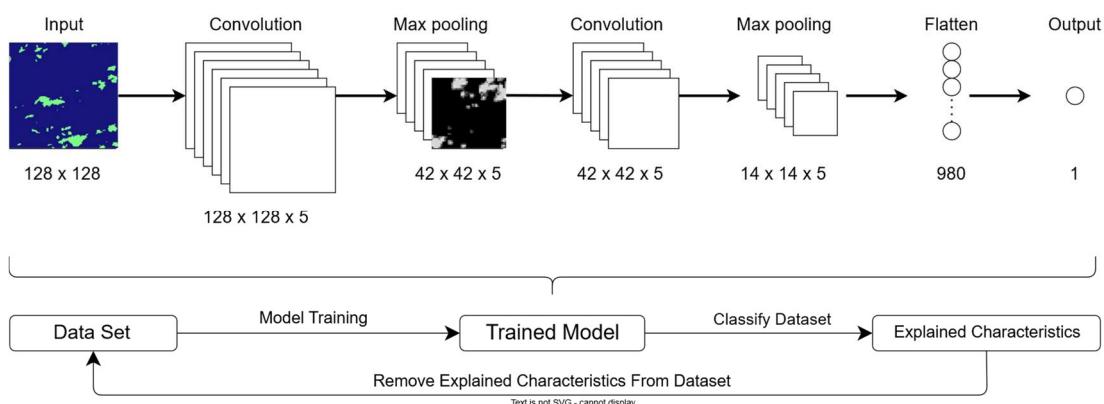
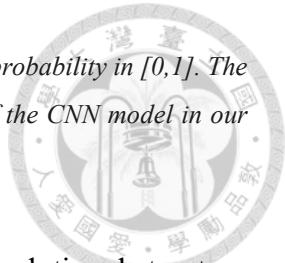


Figure 3-1. A schematic diagram shows the data flow of the CNN model and our framework. Each kernel of the convolution layers is 3×3 , and the activations are ReLUs. The pool size of the max-pooling layers

is 3×3 , and the activation function of the output layer is sigmoid to represent the probability in $[0,1]$. The flowchart shows our framework for retrieving the physics process and the role of the CNN model in our framework.



This typical network can be applied to image classification. The convolutional structure preserves the spatial distribution's characteristics while reducing trainable parameters. Since the logic of identifying the aggregation hides in the connection between layers of CNN, a common way to understand the logic is to visualize the signal output from the hidden layer and connect it with the physics meaning we are familiar with. This study will visualize the first max-pooling layer right after the first convolution layer showed in Figure 3-1 to find out the connection with physical variables.

We intentionally constructed a CNN model with only a few layers. While deeper networks with more trainable parameters might increase classification accuracy, they often hide the physical signals through complex, high-level abstractions, making interpretation difficult. By limiting the network depth, we may ensure that the signals captured in the first max-pooling layer preserve more original spatial information, allowing us to directly trace the learned features back to specific morphological structures of the clouds.

3.2.2 Dataset

To obtain data covering aggregated states, we use outputs from the Vector Vorticity Equation Cloud-Resolving Model (VVM, Jung and Arakawa (2008); Wu et al. (2019)) simulations (Tsai & Wu, 2017). The dataset includes simulations with varying degrees of environmental moistening. Following Tsai and Wu (2017), we identify distinct modes in cloud size distribution to label "aggregated" versus "non-aggregated" cases. Our framework treats the cloud water field as the primary input for the CNN to classify these states. Specifically, we use the Cloud Water Path (CWP) in each column ($CWP_{i,j}$), as our

input variables, which is calculated by

$$CWP_{i,j} = \sum_k \rho_k \cdot (qc + qi)_{i,j,k} \cdot 10^6 \cdot \Delta z_k$$

where $(qc + qi)_{i,j,k}$ represents the mixing ratio of the cloud condensate ($kg \cdot kg^{-1}$) in each grid point. ρ_k represents the density in the specific height, and Δz_k represents the grid height in the stretching grid.



The simulations are initialized with profiles based on the GATE Phase III field campaign. The imposed large-scale forcing is prescribed as constant in time but modulated to create various environmental conditions. Specifically, the large-scale moistening effect is controlled by multiplying a scale factor ranging from +6 (very moist) to -6 (very dry) with an interval of 1.5. This results in nine distinct experiments representing a spectrum of environments. Following Tsai and Wu (2017), we identified that the simulations with the strongest moistening factors (+6 and +4.5) successfully developed into aggregated convection, while the drier cases (factors -4.5 and -6) remained in a scattered, non-aggregated state.

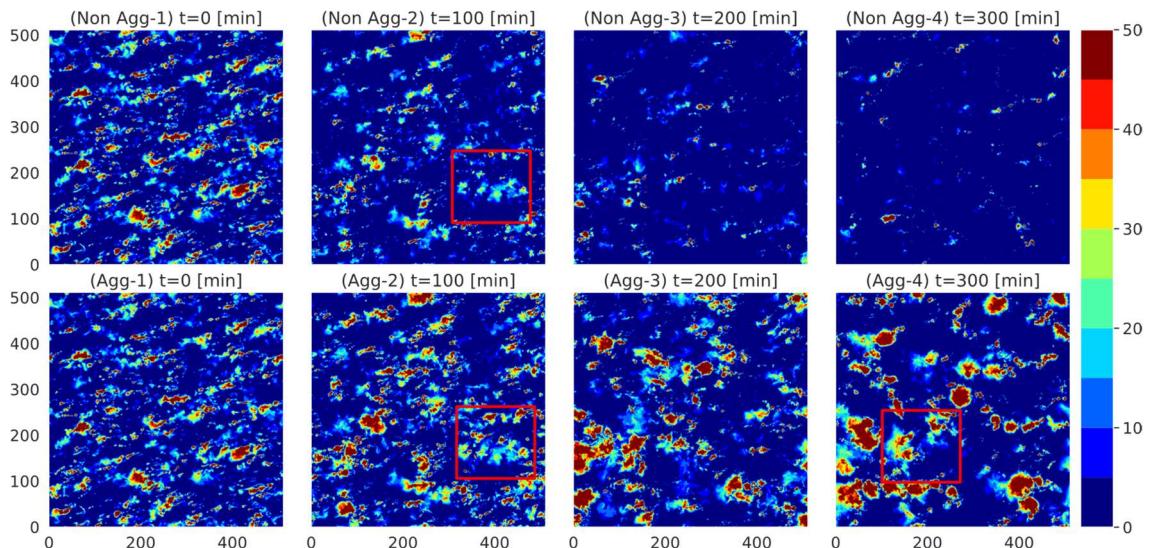


Figure 3-2. (Non agg-1) to (non agg-4) shows the examples of the nonaggregated case, and (Agg-1) to (Agg-4) shows the aggregated case with the cloud water path values in ($kg \cdot m^{-2}$). Each subgraph shows the simulations' 512×512 (km) domain. The solid red lines represent the area that may be difficult to

classify visually. Those areas can be the clouds before elimination in the nonaggregation cases, and the clouds are not at the core of the convection in the aggregation cases.

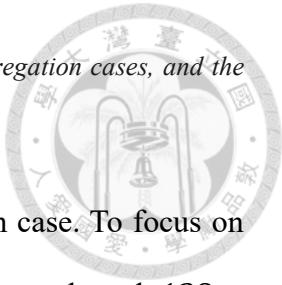


Figure 3-2 shows snapshots of an aggregation and a non-aggregation case. To focus on the characteristics of aggregated convection and its environment, we selected $128 \times 128 \text{ km}^2$ cloud condensate path crops. The selection of the crop size is motivated by the characteristic length scale of convective aggregation. While individual convective cells typically operate on scales of a few kilometers, the organization process behave as the spatial separation of moist and dry regions on the mesoscale (tens to hundreds of kilometers). A 128 km domain is sufficiently large to capture the morphological clusters of aggregated convection and their surrounding dry area, yet small enough to serve as a local sample for the neural network, allowing us to augment the dataset and focus on the local texture of organization rather than global statistics.

3.2.3 Framework: Iterative Feature Removal

A CNN model with many parameters trained on a limited dataset can suffer from overfitting or rely on trivial features. To address this and ensure explainability, our framework is designed to iteratively modify the training dataset. The logic is as follows: We train the model, identify the leading characteristic the model relies on (e.g., average CWP), and then remove or mask that characteristic from the dataset. We then re-train the model. If the model can still predict accurately, it must be using a different set of features.

It is important to clarify the mechanism of "Feature Removal" in this framework. We do not remove the input variable (Cloud Water Path) entirely, as it serves as the carrier of morphological information. Instead, we perform Information Masking on the spatial field

to eliminate specific statistical properties that are known to be highly correlated with aggregation. For example:

1. Masking Intensity: By converting the continuous CWP field into a binary mask (0 or 1), we remove the functional dependence on cloud thickness/intensity while preserving the spatial layout.
2. Masking Coverage: By restricting the dataset to samples with similar cloud fractions, we remove the statistical dependence on cloud coverage area.

This process is analogous to causal discovery via intervention. By systematically blocking the "easy" pathways (intensity and coverage) that the neural network would typically use for classification, we force the model to seek alternative discriminative features hidden in the remaining information, specifically, the spatial distribution and edge complexity (morphology). This ensures that the learned features represent nonlinear structural drivers rather than simple scalar correlations.

By repeating this process in this framework, we may dig out the deeper, nonlinear morphological relationships.

3.3 Iterative Extraction of Morphological Features

Iteration 1: The Role of Average Cloud Water Path

In the first round, the trained model based on the complete CWP information achieves an accuracy of about 98%. By analyzing the predicted probability against the domain-averaged CWP (Figure 3-3A), we find that cases with averaged CWP greater than $1 \text{ kg} \cdot \text{m}^{-2}$ are easily classified as aggregated. This is a trivial feature for distinguishing between aggregated and non-aggregated cases according to our knowledge of convective aggregation. To force the model to learn more detailed variables, we removed these high-

average-CWP cases from the dataset.

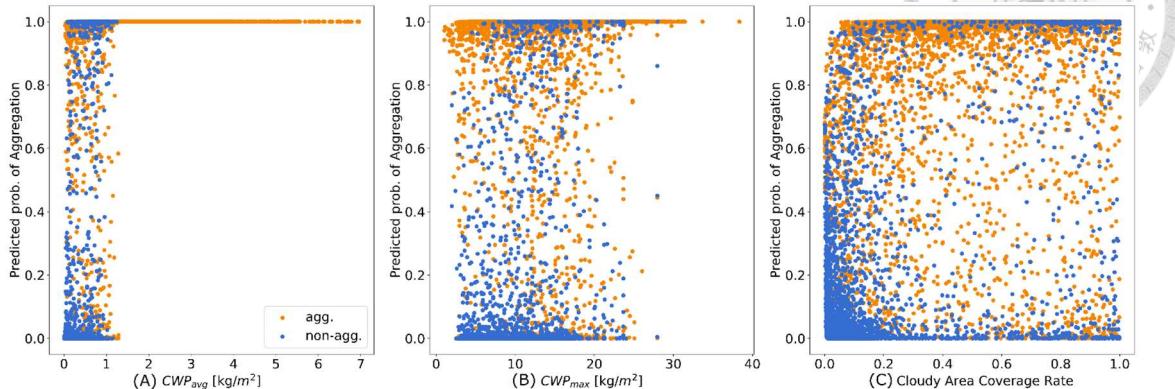


Figure 3-3. The model predicted probabilities of the aggregation (y-axis in each subgraph) from the CNN model in the different iterations and the potential property derived from the CWP values (x-axis in each subgraph). Blue points in each subgraph stand for the actual label of the training data are nonaggregation cases, and the orange points stand for the aggregation cases. (a) The 128×128 km 2 domain averaged CWP ($\text{kg} \cdot \text{m}^{-2}$), (b) the maximum of the CWP ($\text{kg} \cdot \text{m}^{-2}$) in the 128×128 km 2 crops. (c) The cloudy area covers 128×128 km 2 crops after being masked by the value of $2 \text{ kg} \cdot \text{m}^{-2}$.

Iteration 2: The Role of Maximum Intensity

The model trained on the remaining dataset still achieved 95% accuracy. Figure 3-3B shows that the maximum CWP can be a potential discriminator—aggregated cases tend to have higher extrema. To eliminate this intensity-based feature and focus on the spatial distribution, we masked the CWP information by setting a threshold ($2 \text{ kg} \cdot \text{m}^{-2}$). Grids with values above this are set to 1, and others to 0. This retains the *cloud shape* (morphology) while removing the *intensity distribution*.

Iteration 3: The Role of Cloud Coverage

In the third round, even with intensity masked, the model persists with 90% accuracy. We hypothesized that cloud coverage rate (Cloud Fraction) is the key. As shown in Figure 3-3C, low coverage rates are associated with non-aggregation. To challenge the CNN model further, we limited the dataset to cases with cloud coverage between 5% and 10%, which are the hard cases in Figure 3-3C

After these iterations, the modified dataset is difficult even for humans to distinguish visually (Figure 3-4A, B). Yet, the CNN model still achieves 73% accuracy (better than random guessing). This suggests the existence of a hidden morphological relationship that is related to the mechanisms of cloud organization.

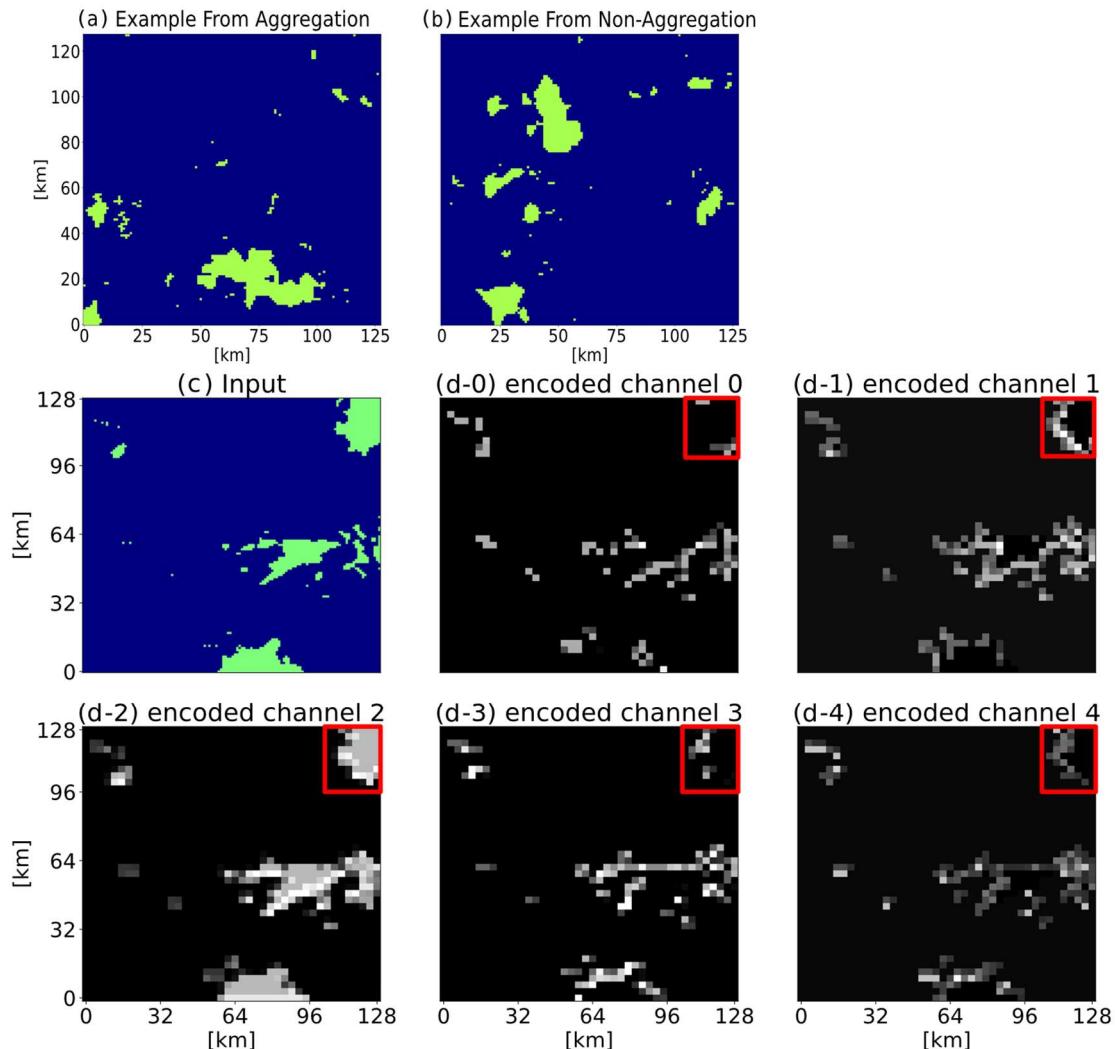


Figure 3-4. (a) and (b) show the example of the comparison between nonaggregation and aggregation labeled data in the training dataset. The cropped cloud water path field on a 128×128 km 2 is masked by a $2 \text{ kg}\cdot\text{m}^{-2}$ threshold. The green grid will be regarded as one, while dark blue grids will be regarded as zero. The ones and zeros will be inputted into the input layer of the CNN model. The graph in (c) shows the referenced cloud water (masked) inputted into the CNN models. From (d-0) to (d-4), show the graphs of the encoded channel 0 to 4, representing the max-pooling layer with the 42×42 grid or the 128×128 km domain. Each grid concludes the signal from the 3×3 km 2 box. The brightness in the graphs shows the strength of the signal outputted. Each channel has its scale, so we will only focus on the relative strengths

and the spatial distributions. The channel order may vary due to the randomness of the training process; however, the characteristics of the relative strengths and the spatial distributions will remain.



3.4 Discussion: Decoding the Hidden Layer

To quantify the complexity of the cloud edges, we calculated the fractal dimension of the cloud clusters. In this study, the fractal dimension is applied to the binary cloud mask (where $CWP > 2 \text{ kg/m}^2$). A higher fractal dimension indicates a more complex, rougher boundary, whereas simple, smooth shapes typically yield lower dimensions.

To understand the "black box" logic of the final iteration, we visualize the signals from the CNN's first max-pooling layer with 5 channels (Figure 3-4D, channel 0 to 4). Strong signals appear at the center of clouds in one channel (Channel 2) and at the edges in others (Channels 0, 1, 3, 4).

This indicates that the model is looking at the fractal dimension or the complexity of the cloud boundaries. To quantify this, we analyzed the correlation (R^2) between the signal sums of specific channels and physical properties (Cloud Fraction vs. Fractal Dimension).

	Cloud Fraction		Fractal Dimension		R^2 values
	Non-Aggregation	Aggregation	Non-Aggregation	Aggregation	
Channel 0	0.33	0.09	0.17	0.66	1.0
Channel 1	0.49	0.11	0.09	0.61	0.9
Channel 2	0.88	0.63	0.01	0.16	0.8
Channel 3	0.51	0.12	0.07	0.62	0.7
Channel 4	0.26	0.08	0.23	0.68	0.6
					0.5

Table 3-1. R^2 values of the cloud fraction/fractal dimension and the channel signal sums corresponded to the model used in Figure 3-4 (d-0) to (d-4), separated by the aggregated and non-aggregated cases.

Table 3-1 reveals a striking result: In non-aggregated cases, the CNN signals correlate

strongly with Cloud Fraction (Channel 2, $R^2 = 0.88$). This means the model may simply count how many cloud grids there are to confirm the non-aggregation cases. In aggregated cases, the CNN signals correlate strongly with Fractal Dimension (Channels 0, 1, 3, 4, $R^2 > 0.6$). This implies that when convection is aggregating, the geometry of the cloud edges and their nonlinear complexity become the dominant feature for classification. This confirms that the CNN has learned that aggregated convection is morphologically distinct not just in size, but in the nonlinearity of its spatial structure.

The result highlights a fundamental difference in the underlying logic of convective states. The high correlation in non-aggregated cases suggests that the model finds a linear relationship with a simple scalar metric—cloud fraction. In contrast, the difference in linearities observed in aggregated cases indicates that scalar metrics are insufficient. The model instead shifts its focus to the fractal dimension, which represents the nonlinearity in the geometry of the cloud structure. This confirms that as convection aggregates, its defining characteristic transitions from “how much cloud is there” (linear) to “how complex is the cloud shape” (nonlinear morphology).

Also, the choice of using CWP as the input variable is strategically aligned with observational capabilities. The processed dataset, after removing trivial intensity features, may be similar to the imagery obtained from satellite liquid water or ice channels. Unlike nadir-view visual images that may be obscured by cirrus shields, these channels provide a more direct representation of the convective core structure. Thus, demonstrating that the CNN can identify aggregation solely from the morphology of CWP suggests that this framework has the potential to be adapted for diagnosing convective organization from

geostationary satellite observations.



3.5 Summary

In this chapter, we introduced a data-driven framework to retrieve essential cloud characteristics using CNN models and VVM simulations. By employing an iterative feature removal strategy, we analyze the layers of the "black box." We identified that while average CWP and cloud coverage are primary indicators, the nonlinear morphological features—specifically the cloud coverage and edge complexity of clouds—can be the hidden features that distinguish aggregated convection when other statistical features are ambiguous.

This chapter demonstrated that this framework of deep learning can successfully extract morphological features from mesoscale fields to classify convective organization. However, in the context of Global Storm-Resolving Models (GSRMs), many crucial morphological processes occur at even finer scales—specifically, the subgrid-scale of GSRMs.

One such critical phenomenon that related to convective organization is the Cold Pool, whose dynamic edges (the morphology of which drives new convection) are often smaller than the grid size of current global models. While Chapter 3 focused on identifying existing cloud patterns, Chapter 4 will shift the focus to reconstructing these subgrid details. We will move from a classification task (CNN) to a segmentation and downscaling task using a U-Net architecture, aiming to predict the high-resolution morphology of cold pools from coarse-grained environmental fields, thereby providing a pathway for future parameterization.

Chapter 4



Capturing Subgrid-Scale Cold Pool Morphology: A U-Net Based Morphological Reconstruction

4.1 Introduction

In the previous chapters, we demonstrated how deep learning can identify morphological features at the large scale (SSW evolution via VAE) and mesoscale (convective aggregation via CNN). In this chapter, we descend to the subgrid scale—an important scale for the development of the next generation of climate modeling.

The interaction between convective clouds, driven by cold pool dynamics, plays an important role in triggering convective systems and producing extreme rainfall events (Khairoutdinov & Randall, 2006; Moseley et al., 2016; Tompkins, 2001). Observations and simulations show that the morphology of cold pools—specifically their intensity gradients and spreading boundaries—strongly influences nearby convective initiation (Böing, 2016; Feng et al., 2015; Haerter, 2019; Haerter et al., 2019; Khairoutdinov & Randall, 2006). To better capture the role of cloud convection in climate, recent development of Global Storm-Resolving Models (GSRMs, Stevens et al. (2019)) has raised the focus on how variability in convective systems interacts with large-scale circulation.

However, the horizontal resolution currently used in GSRMs (typically a few kilometers)

may be insufficient to fully resolve the physical processes related to cold pools. Since the dynamics of cold pools are highly nonlinear (Hohenegger & Schlemmer, 2014; Torri & Kuang, 2019), accurately capturing these processes requires models capable of representing complex, nonlinear physical interactions with high spatial simulation resolution (Kao et al., 2025; Moseley et al., 2016; Wang et al., 2024).

In the context of this dissertation, we propose that deep learning can serve as a bridge across this resolution gap. We employ high-resolution Large-Eddy Simulation (LES) to generate the "ground truth" morphology of cold pools. We then develop a U-Net deep learning model—a network designed for image segmentation and reconstruction—to predict the high-resolution distribution of cold pools using coarsened physical variables. This allows us to assess whether the nonlinear morphological features of cold pools can be retrieved from variables resolved by GSRMs, providing insights for future subgrid parameterizations.

4.2 Methodology

4.2.1 Data Description: Generating Ground Truth Morphology

The dataset is generated using the Vector Vorticity Equation Cloud-Resolving Model (VVM) to simulate the evolution of cold pools. This model has been run at 100 m horizontal resolution to capture the fine-scale structures (Kao et al., 2025; Tsai & Wu, 2016; Wang et al., 2024; Wu & Arakawa, 2011). We adopt simulations from Kao et al. (2025), which represent typical oceanic conditions and encompass a variety of convective regimes.

To ensure the physical realism of the cold pool structures in our training data, the simulations incorporate comprehensive physical parameterizations. Specifically, the model uses the RRTMG scheme (Iacono et al., 2008) for radiation processes, the P3 scheme (Morrison & Milbrandt, 2015) for microphysics, and a first-order closure for turbulence representation (Shutts & Gray, 2006). This setup ensures that the simulated cold pools are driven by realistic interactions between radiative cooling, precipitation evaporation, and boundary layer mixing.

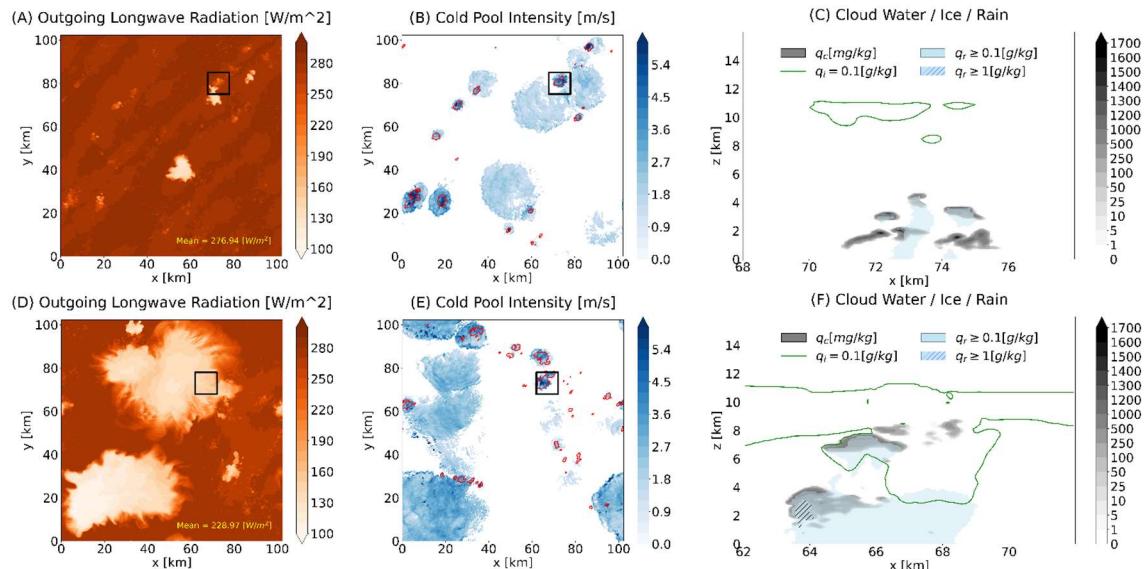


Figure 4-1. Cold pool and convection characteristics from Large-Eddy simulations. Snapshots from high-resolution Large-Eddy Simulation (LES) illustrating two convective cases, case 1 from (A) to (C), and case 2 from (D) to (F). Outgoing Longwave Radiation (OLR, as shown in A and D) indicates cloud-top temperatures and convective activity with the domain-averaged value on the bottom right. Cold pool intensity (as shown in B and E) computed by vertically integrating negative buoyancy over the boundary layer, and red contours denote precipitation rates exceeding 1 mm h^{-1} . Vertical cross-sections (as shown in C and F) show the vertical structures of cloud water, ice, and rainwater within the boxed regions in (A) (B) and (D) (E). The top row from (A) to (C) represents shallow convection with limited OLR depression and small cold pools, while the bottom row from (D) to (F) shows deep convection with extensive anvil clouds and deep vertical development.

Figure 4-1 demonstrates the variability of cold pools, in which cold pool intensity B is calculated by the vertically integrated negative buoyancy over the boundary layer,

following (Feng et al., 2015; Rotunno et al., 1988; Tompkins, 2001).

$$B = \sqrt{2 \times \int_0^h -g \frac{\theta_p - \bar{\theta}_p}{\bar{\theta}_p} dz} \text{ [m/s]}$$



where θ_p and $\bar{\theta}_p$ specify the cold pool and horizontal averaged virtual potential temperature, h specifies the height when negative buoyancy is greater than $-0.003 \text{ [m} \cdot \text{s}^{-2}\text{]}$ as used in Feng et al. (2015). Importantly, cold pools do not always coincide spatially with precipitation (Figure 4-1B, 4-1E), emphasizing that their morphology is driven by complex boundary-layer processes that cannot be inferred solely from rainfall.

The selection of input variables, horizontal winds (u, v), potential temperature (θ), and relative humidity (RH), presents our target physics of cold pool. While the cold pool is fundamentally a thermodynamic anomaly (dense, cold air), its propagation and ability to trigger secondary convection are dynamically driven by the leading gust front. It is important to note that simply interpolating these high-resolution dynamics from a sub-kilometer scale to a coarse grid (e.g., via linear interpolation) results in a significant loss of physical accuracy. The sharp gradients characterizing the gust front are smoothed out in the coarsening process, erasing the signal of mechanical lifting. Therefore, our deep learning task is not only downscaling, but also involving the nonlinear reconstruction of these sharp morphological boundaries that are dynamically consistent with the coarse-grained environmental kinematic fields (specifically the divergence signatures in u and v), which traditional interpolation methods fail to preserve.

4.2.2 U-Net Architecture for Morphological Reconstruction

Unlike the CNN used in Chapter 3 which compressed spatial information into a scalar probability, here we adopt the U-Net (Ronneberger et al., 2015) deep learning model, which is a variation of CNN. The U-Net architecture preserves spatial information through skip connections, making it ideal for reconstructing high-resolution features from low-resolution inputs.

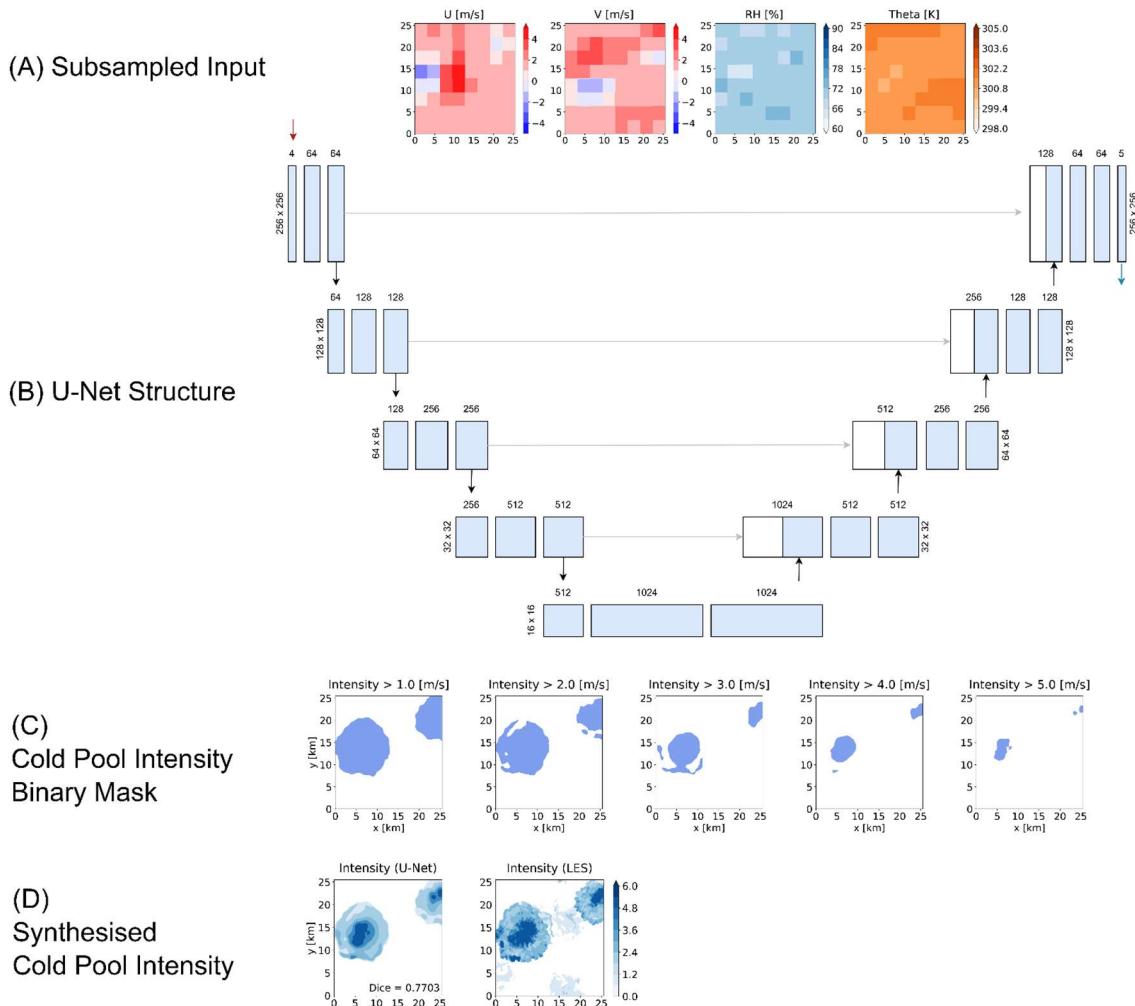


Figure 4-2. U-Net framework for cold pool intensity prediction. This figure presents the data processing pipeline and U-Net architecture used to predict high-resolution cold pool intensity from coarse-resolution atmospheric fields. (A) Subsampled Input: Low-resolution input fields (wind, relative humidity, potential temperature) at 0.8, 1.6, 3.2, 6.4 km resolution; we take 3.2 km resolution as an example in this figure. (B) U-Net structure: The deep learning network structures used in this study, which predict high-resolution cold pool intensity from low-resolution physics fields. (C) Cold pool intensity binary masks: The model outputs five binary masks corresponding to cold pool intensity thresholds from 1 to 5 m/s. These are

compared with LES ground truth to compute Dice loss and guide model training. (D) Synthesized cold pool intensity and LES ground truth: The binary masks are combined to reconstruct the predicted cold pool field (left), which can be compared visually against the LES ground truth (right). The Dice score shown in the bottom right quantifies the prediction quality.



Our strategy involves taking the high-resolution model outputs and coarsening them to various resolutions (0.8, 1.6, 3.2, and 6.4 km) to serve as U-Net inputs. The model is trained to predict the high-resolution (100 m) cold pool intensity. Instead of regressing the continuous cold pool intensity directly, we discretized the field into five binary masks at intervals of 1 m/s (Figure 4-2C). This design fits the physical nature of cold pools, which function as density currents characterized by sharp gradients (gust fronts) rather than smooth variations. By training the U-Net to predict these discrete intensity contours, we force the model to explicitly learn the morphological boundaries and the internal structure of the cold pools. This avoids the common issue in regression where models produce smoothed, blurry fields that fail to capture the sharp dynamic edges essential for triggering secondary convection.

4.2.3 Data Balancing

To ensure the model generalizes well across different convective regimes, we applied a class balancing strategy. The dataset was divided into six predefined intervals based on the cold pool fraction to ensure a balanced representation of both strong and weak cold pool events in the training process. The final dataset consists of 3,600 samples, split into 2,880 samples for training and 720 samples for testing. Also, to evaluate the U-Net's capability to bridge the resolution gap, we train four models with identical architectures using inputs at different coarse resolutions.

4.3 Results: Reconstructing Sub-grid Morphology

4.3.1 Overall Performance



Figure 4-3 compares the U-Net (C-F) and a baseline Linear Regression (LR) model (G-J). We define cold pool boundaries ("Edges", Figure 4-3B) to assess how well the model captures the spreading morphology.

Specifically, the cold pool edges are determined by extrapolating the current cold pool position based on a propagation speed derived from the density current theory. This extrapolation estimates the potential spreading area of the cold pool over a 10-minute interval, serving as an estimation for assessing the model's ability to capture the morphology in all the intensity prediction. The orange contours in Figure 4-3 indicate this "Edges".

To benchmark the performance of the U-Net, a pixel-wise Linear Regression (LR) model was employed. It is acknowledged that LR represents a relatively simple baseline compared to more advanced regression methods. However, the primary purpose of using LR here is to visualize and quantify the nonlinearity of the subgrid reconstruction problem. If the relationship between the coarse-grained environment and fine-scale cold pool morphology were linear, the LR model would yield reasonable approximations.

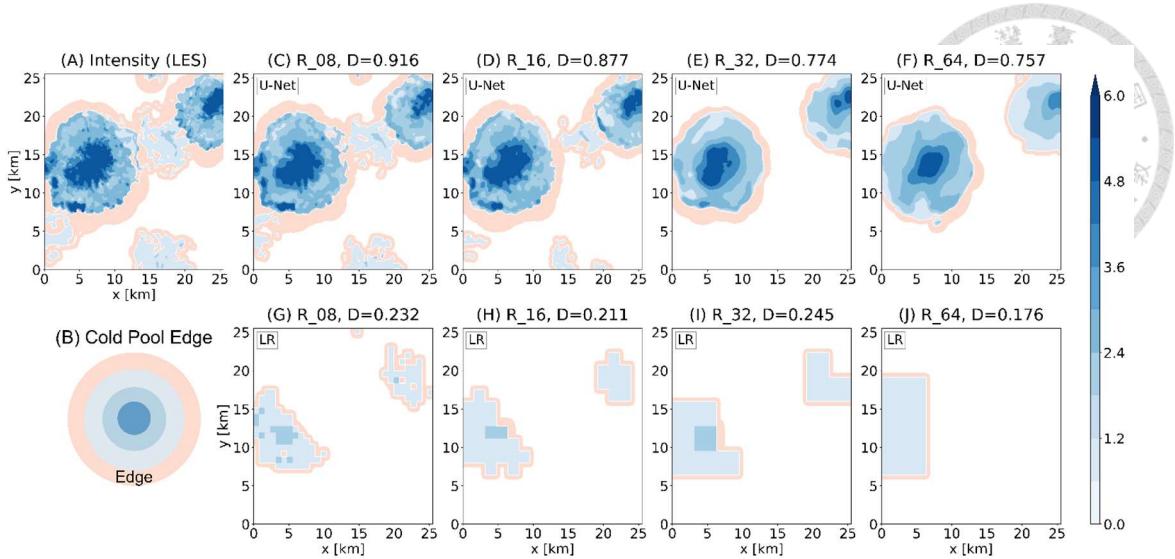
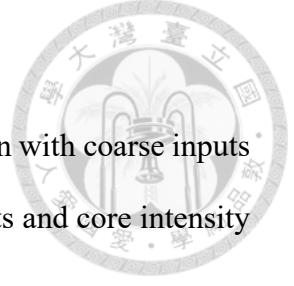


Figure 4-3. Comparison of cold pool predictions using U-Net and linear regression models. (A) High-resolution cold pool intensity from LES (100 m resolution). (B) Cold pool boundaries (marked as “Edge” in B and with orange boundary areas in C to J) based on 10-min propagation. (C to F) U-Net predictions at input resolutions of 0.8, 1.6, 3.2, and 6.4 km, and (G) to (J) Linear regression model predictions at the same resolutions. The value D on each subgraph’s title represents Dice coefficients, which quantify the prediction accuracy.

At finer resolutions (0.8 and 1.6 km), U-Net accurately captures both the spatial extent and intensity of cold pools. As the resolution becomes coarser (3.2 and 6.4 km), the prediction quality naturally declines, showing reduced heterogeneity. However, the contrast with the LR model is noticeable. The obvious contrast in performance between LR (Figure 4-3 G-J) and U-Net (Figure 4-3 C-F), where LR produces fragmented, blocky prediction while U-Net successfully reconstructs coherent boundaries, serves as a demonstration that the mapping from environmental variables to subgrid morphology is highly nonlinear. The LR thus establishes a necessary performance lower bound, confirming that the structural complexity of cold pools requires the hierarchical feature extraction capabilities of deep neural networks.



4.3.2 Detailed Structural Analysis

Figure 4-4 provides a detailed look at the 3.2 km resolution case. Even with coarse inputs (Figure 4C-F), the U-Net successfully reconstructs the sharp gradients and core intensity of the cold pools.

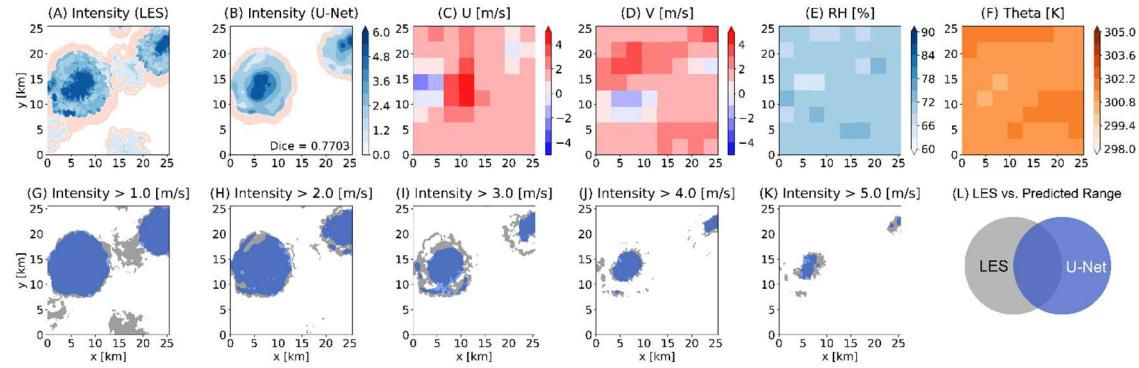


Figure 4-4. Evaluation of U-Net cold pool predictions at 3.2 km resolution. This figure provides a detailed comparison between the cold pool intensity from LES and the U-Net predictions. Figure (A) represents the example of cold pool intensity from LES at 100 m resolution, and (B) the corresponding synthesized U-Net prediction. The corresponding input low resolution physics fields are shown from (C) to (F), we take 3.2 km resolution input fields in this example. From (G) to (K) compare the ground truth cold pool intensity masks (gray regions) with U-Net predicted intensity masks (blue regions) at different intensity thresholds (1.0, 2.0, 3.0, 4.0, and 5.0 m/s). (L) presents a schematic of the intensity mask ranges predicted by LES and U-Net.

While minor differences appear near the boundaries (weak intensity regions), the model captures the core structure (intensity > 3.0 m/s) effectively.

4.3.3 Statistical Validation

We also statistically compare the performance across all resolutions. Figure 4-5 show that the U-Net consistently outperforms the LR model across all resolutions, confirming its robustness in handling nonlinear spatial patterns.

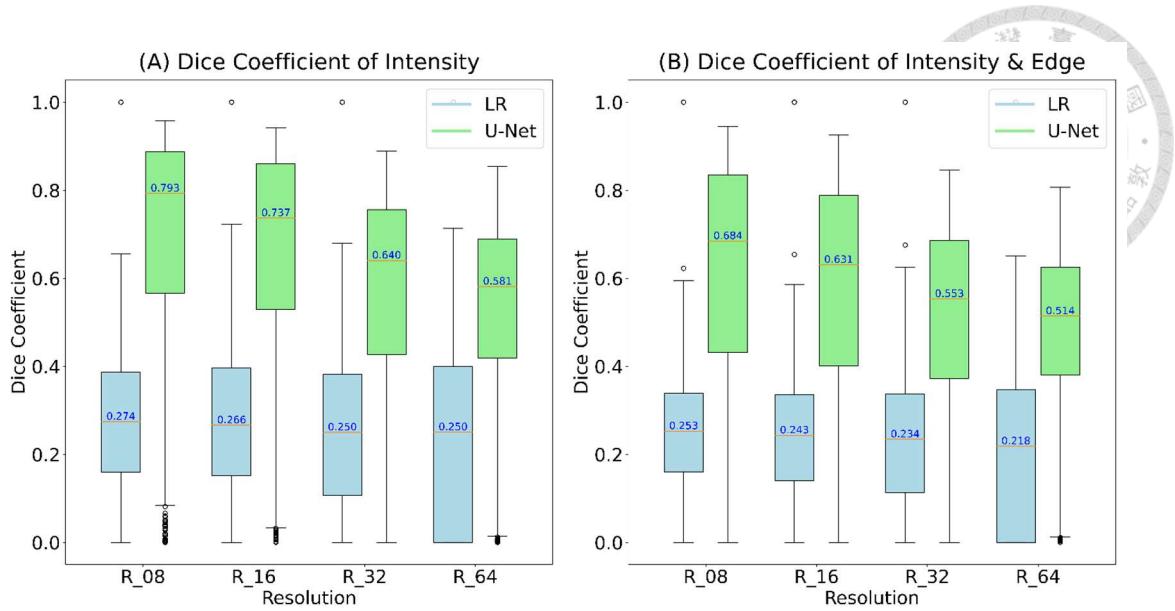


Figure 4-5. Performance comparison of U-Net and linear regression in predicting cold pools. Boxplots show Dice coefficients for U-Net (green) and linear regression (blue) at four resolutions (R_08: 0.8, R_16: 1.6, R_32: 3.2, R_64: 6.4 km) using subsampled input. At each resolution, the Dice coefficients are computed based on comparisons over 720 test samples. (A) represents the cold pool intensity prediction, while (B) covers the cold pool intensity and boundary prediction. The result shows that U-Net model outperforms linear regression model across all resolutions.

4.4 Explainability: Identifying Key Morphological Features

To interpret the U-Net model and ensure it is learning physically meaningful features, we conduct a sensitivity analysis by selectively masking input fields. This aligns with the goal of the dissertation.

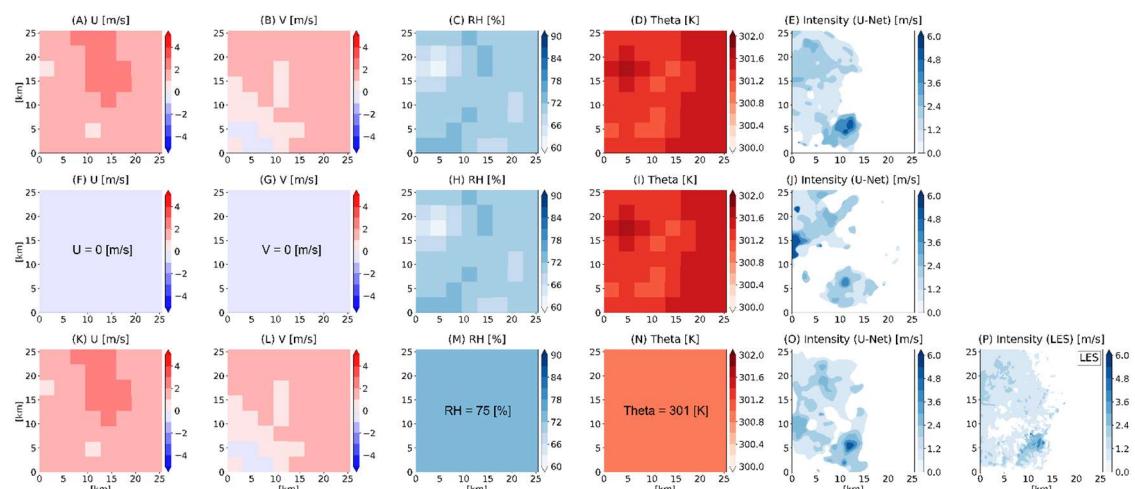


Figure 4-6. Sensitivity analysis of cold pool prediction using different input fields. Each row shows U-Net

input maps (U, V, RH, θ) on 3.2 km resolution (A–D, F–I, and K–N), with the corresponding cold pool intensity from U-Net prediction (E, J, and O) and LES (P). Group 1 from (A) to (E) includes all variables included, group 2 from (F) to (J) masked the dynamic variables (the wind components), and group 3 from (K) to (O) mask the thermodynamic variables (the RH and θ). The ground truth of cold pool intensity from LES is shown in (P). Results show that wind information may be more critical for U-Net cold pool prediction.

Comparing Figure 4-6 from F to J (dynamic masked) and Figure 4-6, from K to O (thermodynamic masked) reveals an insight: Wind fields seem to be more important than thermodynamic fields. When wind information is removed, the model fails to locate the cold pools accurately. This suggests that the U-Net may retrieve the divergence pattern as the primary predictor for cold pool structure, while thermodynamic fields primarily refine the intensity distribution.

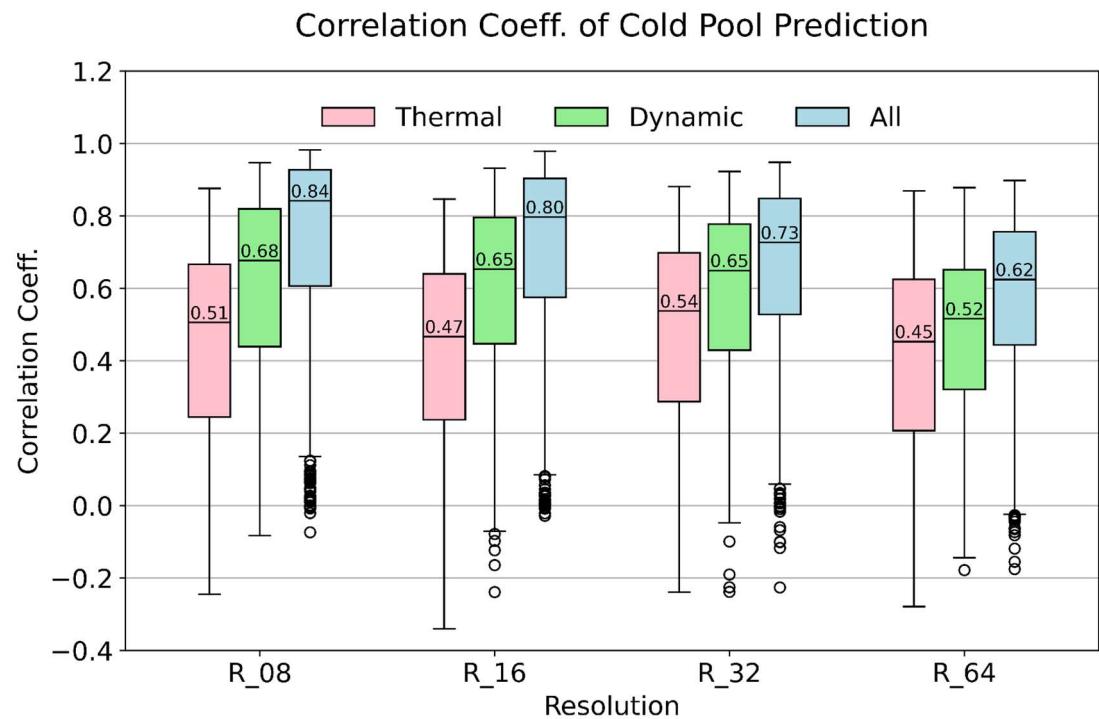


Figure 4-7. Impact of input variables on cold pool prediction performance. Boxplots show correlation coefficients between predicted and ground-truth cold pool fields for four resolutions (R_08 to R_64). Colors represent the combination of input fields: Thermal fields only (pink), dynamic fields only (green), and all variables (blue). The group with the wind fields appears to improve the predictions, especially at finer resolutions. For each resolution and each input group, the correlation coefficients are computed based on 720 test samples.

Figure 4-7 quantifies this finding. The "Dynamic-only" (green) performance is much closer to "All" (blue) than "Thermal-only" (pink) is, especially at finer resolutions. This indicates that near-surface wind may represent a morphological indicator of subgrid cold pool activity.

The sensitivity analysis reveals a crucial physical insight: the dynamic fields (u, v) play a more significant role in morphological reconstruction than the thermodynamic fields (θ, RH). This is particularly relevant for strong convective systems, such as those frequently observed in oceanic deep convection regions. In these regimes, the mechanical lifting at the gust front driven by the propagation of the density current is the primary trigger for new convection. While the cold pool is thermally defined, the thermal characteristics may lag behind the leading edge. In contrast, the kinematic characteristics (surface divergence and convergence) provides an immediate and sharp indicator of the cold pool distribution, especially the boundary location. The U-Net's reliance on wind fields suggests that the model may learn to identify the dynamical representation of the cold pool, specifically the gust front convergence.

4.5 Summary and Discussion

In this chapter, we investigated the capability of deep learning to bridge the scale gap in atmospheric modeling. We developed a U-Net model to predict high-resolution subgrid cold pool morphology from coarse-resolution environmental fields. Our findings reveal that U-Net effectively captures the nonlinear morphological characteristics of cold pools, particularly their edges and intensity gradients, significantly outperforming linear methods. Also, through sensitivity experiments, we identified that the morphology of the

wind field can provide the important information for reconstructing cold pool structures, more than thermodynamic variables alone.



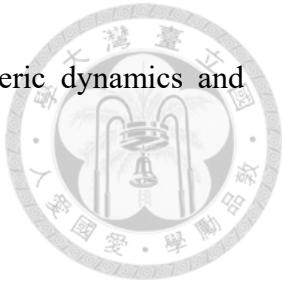
A remaining question is how can the predicted cold pool morphology be physically integrated into a coarse-resolution model to improve convection? We propose a pathway following the approach of Kao et al. (2025). The U-Net predicted cold pool intensity, which represents the vertically integrated negative buoyancy, can be used to estimate the “nonlocal vertical acceleration.” This dynamically derived term can then be introduced as a source term in the vertical momentum equation of the global model. In this way, the morphological features reconstructed by deep learning may be directly translated into the statistical information for a mechanical lifting force, explicitly representing the subgrid triggering mechanism of cold pools.

This study so far completes the third tier of our framework.

- In Chapter 2, we used VAE to represent large-scale global morphology by representing the Z10 in the duration of SSW.
- In Chapter 3, we used CNNs to extract mesoscale convection features by iteratively removing the recognized physics process.
- In Chapter 4, we have used U-Net to reconstruct subgrid-scale details of cold pool intensity from low resolution physics field.

These three chapters collectively demonstrate that, when designed with explainability in mind, the deep learning model can potentially capture and interpret atmospheric morphology across scales. In the next chapter, we will synthesize these findings to discuss the broader implications of using "Morphology" as a physical indicator and how these

"Glass Box" models can transform our understanding of atmospheric dynamics and parameterization.



Chapter 5



General Discussion

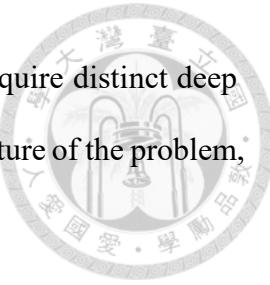
5.1 Linking multi-scale phenomena through morphology

Following the detailed investigations in the previous chapters, we try to address a fundamental challenge in atmospheric science: how to quantify and interpret the complex, nonlinear spatial structures that drive atmospheric dynamics. Traditional statistical methods often reduce these structures to scalar metrics, losing spatial information or distributions. To overcome this, we proposed an explainable deep learning framework that treats "morphology" as a primary physical indicator.

We have applied this framework across three distinct scales: the planetary scale (SSW evolution), the mesoscale (convective aggregation), and the small scale (cold pool dynamics). While the specific physical phenomena differ, the underlying methodological philosophy remains consistent. In this chapter, we synthesize the findings from Chapters 2, 3, and 4 to discuss the role of different neural network architectures in morphological analysis and how we transitioned from "black box" predictions to "glass box" physical interpretations.

5.2 The Role of Neural Network Architectures in Morphological Analysis

A key finding of this research is that different morphological tasks require distinct deep learning architectures. The architecture must align with the physical nature of the problem, which requires the domain knowledge of the problem.



5.2.1 VAE for Representation and Continuity

In Chapter 2, our goal was to understand the evolution of the stratospheric polar vortex.

The physical constraint here is continuity—a vortex does not change abruptly.

- Why VAE worked: The Variational Autoencoder (VAE) is designed to learn a continuous latent probability distribution. By compressing the global morphology into a low-dimensional latent space (z_0, z_1), the VAE naturally enforced a smooth transition between states.
- Comparison: Unlike CNN classification (discrete) or PCA (linear), the VAE captured the trajectory of morphological change, showing that deep learning can serve as a powerful dimensionality reduction tool for global circulation patterns.

5.2.2 Iterative CNN for Causal Feature Extraction

In Chapter 3, the problem was identification: what specific feature makes an environment "aggregated"? The physical constraint here is causality—identifying the driver in the high-dimensional noise.

- Why Iterative CNN worked: A standard CNN is good at feature extraction but may be poor at explanation. By introducing our "Iterative Feature Removal" strategy, we transformed the CNN from a static classifier into a dynamic exploration.
- Comparison: This approach revealed that once trivial features (intensity, coverage) are removed, the model relies on other possible features (fractal dimension). This suggests that we can potentially find the hidden feature through this strategy with

the involvement of our domain knowledge.



5.2.3 U-Net for Reconstruction and Downscaling

In Chapter 4, the challenge was reconstruction: recovering lost details from coarse data. The physical constraint here is spatial correspondence—the output must spatially align with the input environment.

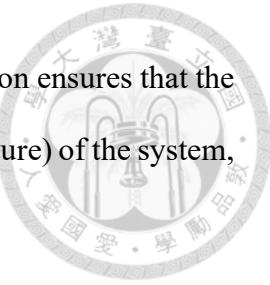
- Why U-Net worked: The U-Net architecture, with its skip connections, preserves spatial information while allowing information to flow between different resolution levels. This allowed us to map coarse-grained environmental conditions to high-resolution phenomena.
- Comparison: This moved beyond simple regression or parameterization schemes, demonstrating that deep learning can effectively generate the subgrid morphology in a physically constrained way, which may be required for low resolution models that need the statistical information from the subgrid.

5.2.4 The Physical Interpretability of Network Operations

To further bridge the gap between deep learning and atmospheric physics, it is important to discuss the implicit physical operations performed by standard neural network components, specifically Normalization and Pooling.

In our framework, normalized input fields (e.g., Z-score standardization) remove the dimensional units and absolute magnitudes, forcing the neural network to learn based on relative spatial variations rather than absolute intensities. For the study of organized convection (Chapter 3), this is particularly important. While the absolute value of cloud water path indicates how much cloud is present, the normalized pattern reveals how the

clouds are arranged. By stripping away the mean intensity, normalization ensures that the model's decision is driven by the nonlinear morphology (shape and texture) of the system, which we identified as a key indicator of the aggregation state.



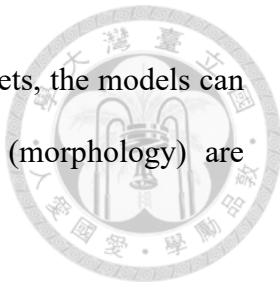
The Max-Pooling layers in our CNN and U-Net architectures serve a function analogous to coarse-graining in physics. By selecting the maximum activation within a window (usually 2×2 pixels), the network effectively filters out small-scale, stochastic noise while preserving the most dominant signals. In the context of convection, this operation emphasizes the convective cores or strong updrafts within a larger system. The combination of pooling layers thus allows the model to hierarchically extract features, moving from local cloud textures to larger mesoscale organizational structures, consistent with the multiscale nature of atmospheric dynamics.

5.2.5 Model Robustness and Sensitivity

To ensure that our "glass box" models rely on physical principles rather than spurious correlations, the sensitivity analyses have been adapted

A common challenge in applying deep learning to extreme events (like SSWs in Chapter 2) is the limitation of sample size. While typical computer vision tasks rely on millions of images, atmospheric phenomena are often rare. Also, the techniques adapted in the computer vision (such as rotation) may not be suitable in the scenario of atmospheric sciences problem. We addressed this by employing data augmentation strategies such as patch-cropping (in Chapter 3) and class-balancing (in Chapter 4). Our results suggest that, with the domain knowledge involved, the morphological features are often more

information-dense than scalar metrics. Thus, even with smaller datasets, the models can achieve some level of reliability if the distinct spatial patterns (morphology) are physically consistent.



5.3 The Choice of Loss Functions

In deep learning, the loss function serves as the guiding force that dictates how a network learns. A critical aspect of our framework is that the choice of loss function is not arbitrary; rather, it represents the translation of domain knowledge into mathematical constraints. In this dissertation, we selected distinct loss functions for each scale to align with the specific morphological characteristics of the atmospheric process being studied.

For the study of Sudden Stratospheric Warmings (Chapter 2), our physical goal was to capture the continuous evolution of the polar vortex. We employed a composite loss function combining Mean Squared Error (MSE) and Kullback-Leibler (KL) Divergence.

- **MSE** acts as a "fidelity constraint," ensuring that the reconstructed geopotential height fields retain their global spatial patterns and intensity.
- **KL Divergence** acts as a "regularization constraint" on the latent space. Physically, atmospheric variability is continuous; the vortex does not teleport from one state to another. The KL term forces the latent distribution to be smooth and continuous, allowing the VAE to represent the SSW evolution as a trajectory rather than discrete classifications. This combination effectively balances spatial accuracy with temporal physical continuity.

In identifying Convective Aggregation (Chapter 3), the objective was to distinguish between aggregation states. Here, we used Binary Cross-Entropy (BCE) Loss. Unlike

MSE, which focuses on pixel-wise reconstruction, BCE operates on probability distributions. It penalizes the model heavily when it is confident but wrong. This forces the CNN to focus on the most discriminative morphological features that separate the aggregated state from the scattered state.

The design choice appeared in the reconstruction of Cold Pools (Chapter 4). Cold pools are physically significant but spatially sparse—they occupy a very small fraction of the total domain compared to the environment. If we had used a standard pixel-wise loss (like MSE or Cross-Entropy), the model could achieve high accuracy by simply predicting "no cold pool" everywhere. To address this class imbalance, we adopted the Dice Loss.

- Physical Implication: Dice Loss measures the overlap between the predicted and actual shapes. It is less sensitive to the large background area and highly sensitive to the successful intersection of the cold pool regions.
- This choice forces the U-Net to focus on the morphology and topology of the cold pool objects (in 5 levels of intensity) rather than the exact intensity value of every background pixel. This aligns with the physical need to understand the area coverage and intensity distribution of subgrid convection.

In conclusion, the selection of loss function reflects our strategy of adapting the "optimization goal" to the "physical nature" of the problem. This demonstrates that defining the loss function is an important step in making deep learning models physically interpretable and robust for atmospheric science applications.

5.4 Morphology in Time

While this dissertation primarily focuses on extracting spatial morphological features

from atmospheric snapshots, an important perspective of deep-learning research in atmospheric sciences is the integration of the time dimension.



In Chapter 2, we touched upon this by mapping the temporal evolution of SSWs into a continuous trajectory within the VAE's latent space. This demonstrates that deep learning can capture the track of morphological change, not just static patterns. However, recent advancements in Autoregressive models and Transformer architectures (e.g., GraphCast, Lam et al. (2023); Pangu-Weather, Bi et al. (2023)) suggest the possibility of learning the governing laws of evolution directly from data. By treating the atmospheric state as a sequence of morphological embeddings, these models may perform data-driven temporal dynamics.

In the context of our framework, this could be achieved by feeding the morphological features extracted by our CNNs or VAEs into a temporal model (such as an LSTM or Transformer). This would allow the system to predict not just what the morphology is (e.g., is it aggregated?), but how it will evolve (e.g., will the cold pool trigger a new cell?), effectively moving from feature extraction to dynamic prediction of 4D atmospheric processes.

5.5 From Black Box to Glass Box: Bridging Data and Physics

The theoretical contribution of this dissertation is the demonstration of interpretability and explainability. The "Black Box" of deep learning has long been a barrier to its acceptance in physics-based sciences. We addressed this by ensuring that "morphology" serves as the common language to balance the prediction score and physical causality.

In Chapter 2, the VAE compressed the stratospheric vortex morphology into a continuous latent space. Physically, the trajectory within this space (z_0, z_1) does not merely represent geometric deformation but serves as a proxy for the dynamical forcing of planetary waves. The continuous transition observed between displacement (Wave-1 dominant) and splitting (Wave-2 dominant) patterns suggests that these events are governed by a fluid spectrum of wave-mean flow interactions rather than discrete regime shifts. The VAE thus provides a "dynamical phase diagram," where the coordinates encode the instantaneous balance of potential vorticity gradients, offering a new metric to quantify the preconditioning and development of stratospheric warming events.

The iterative CNN analysis in Chapter 3 revealed that the complexity of cloud edges (fractal dimension) may be a discriminative feature of aggregation. Physically, this morphological complexity may reflect the thermodynamic interaction between the convective core and the surrounding dry environment. A higher fractal dimension at the boundaries implies a sharp, active interface where moisture gradients are steepening. This may suggest that the aggregation detected by the model corresponds to the convection that has successfully overcome the suppression of dry air entrainment, a critical condition of the self-aggregation

In Chapter 4, the U-Net's ability to reconstruct sharp cold pool boundaries and this may rely more on the near-surface wind fields. This points to a potential physical interpretation: the model is identifying zones of mechanical expanding or lifting. The "edge" of a cold pool is physically the gust front, where density currents collide with the environment to trigger secondary convection. By successfully segmenting these boundaries from coarse-grained data, the deep learning model demonstrates that it captures the feature of low-

level convergence. This implies that the morphology of the cold pool can be a potential proxy of the dynamical lifting mechanism required to trigger or sustain the deep convection, validating the potential of this approach for sub-grid parameterization in global models.

In summary, these approaches validate our hypothesis that morphology serves as a nice bridge between data-driven algorithms and atmospheric physics. By forcing the DL models to reveal their internal logic through these three frameworks, we confirms that the resulting predictions can be physically interpretable. This successful synthesis of deep learning and physical reasoning provides the basis for the summary of findings and future perspectives discussed in the next chapter

Chapter 6



Conclusion and Future Perspectives

6.1 Summary of Major Findings

This dissertation has established a framework for applying explainable deep learning to identify and analyze the nonlinear morphological features of atmospheric processes. The major findings are summarized as follows:

1. Large-Scale Evolution: We demonstrated that a Variational Autoencoder (VAE) can compress the complex evolution of Sudden Stratospheric Warmings (SSWs) into a 2D latent space. The resulting phase diagram provides a more continuous and physically interpretable trajectory of vortex displacement and splitting than traditional PCA methods.
2. Mesoscale Aggregation: Using an Iterative CNN framework, we identified that the transition to convective aggregation is characterized not just by cloud coverage, but by specific morphological features—specifically the fractal dimension of cloud edges. This confirms that the geometry of cloud systems contains hidden information about their convective state, suggesting that the morphological framework could potentially be extended to identify other forms of organized convection in the future.
3. Subgrid-Scale Reconstruction: We showed that a U-Net architecture can accurately reconstruct high-resolution cold pool morphology from coarse-grained environmental fields. Sensitivity analysis revealed that the near-surface wind field is an important predictor, providing a roadmap for the development of future subgrid

parameterizations in global models.



6.2 Contributions

This research contributes to the field of atmospheric science in two ways:

- **Methodological Perspective:** We provided a set of strategies (Latent Space Mapping, Iterative Feature Removal, Resolution-Aware Segmentation) for opening the "black box" of deep learning. These methods allow researchers to use DL not just for prediction, but for physical discovery.
- **Theoretical Insight:** We reinforced the importance of morphology as a quantifiable physical variable. We showed that "shape" and "pattern" can be connected to the nonlinear dynamics of the atmosphere. This connection is beyond atmospheric scales from the stratosphere to the boundary layer.

6.3 Future Perspectives

Building on the foundation laid by this dissertation, several promising directions for future research emerge.

6.3.1 Investigating Cold Pools over Complex Topography

While Chapter 4 focused on idealized oceanic conditions, real-world convection often interacts with complex terrain. In regions like Taiwan, the interaction between cold pools and steep topography is an unavoidable challenge, and the mechanism for triggering secondary convection and extreme rainfall requires further study to be deeply understood.

Future work will apply the deep learning framework to high-resolution TaiwanVVM simulations (P.-Y. Chen & C.-M. Wu, 2025). The challenge here will be to incorporate topographic information into the deep learning model. We hypothesize that the morphology of cold pools over terrain will reveal some distinct nonlinear features (e.g., blocked flow, channelization) that traditional parameterizations struggle to capture, but which a morphology-aware deep learning model could resolve.

6.3.2 Integration into Global Models (Online Implementation)

One of the goals of the offline training conducted in Chapter 4 is online implementation. Future studies could involve embedding the trained U-Net directly into a GSRM as a neural network parameterization. This would allow the global model to generate the statistical information from sub-grid cold pool effects at each time step, potentially improving the representation of the convection and precipitation extremes.

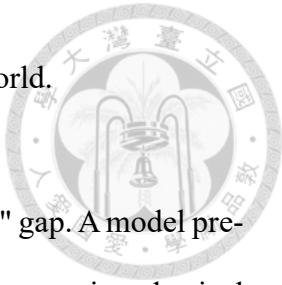
6.3.3 Physics-Informed Deep Learning

To further ensure the robustness discussed in Chapter 1, future models could incorporate physical constraints directly into the loss function (Physics-Informed Neural Networks, PINNs). For instance, ensuring that the reconstructed cold pool mass satisfies continuity equations. Combining the morphological strengths of CNNs with the strict constraints of physics-based equations represents the next steps in this research.

6.3.4 Transfer Learning

One of the most promising applications of the framework established in this dissertation is the transferability to observational data. The morphological features identified in simulations—such as the fractal dimension of aggregated clouds or the sharp gradients of

cold pools—are consistent physical signatures that exist in the real world.



Future work can employ Transfer Learning to bridge the "Sim-to-Real" gap. A model pre-trained on high-resolution simulations (e.g., VVM or other LES) to recognize physical morphology can be fine-tuned using smaller datasets of satellite imagery or Doppler radar reflectivity. This would allow us to apply the "morphology as a physical indicator" approach to real-time monitoring of convective organization or sub-grid severe weather precursors that are currently difficult to parameterize in operational models.

In conclusion, this dissertation demonstrates that by viewing the atmospheric sciences problem through the perspective of explainable deep learning, the "morphology" can potentially serve as a powerful key to revealing the nonlinear physics process.

References



Agrawal, S., Barrington, L., Bromberg, C., Burge, J., Gazen, C., & Hickey, J. (2019). Machine Learning for Precipitation Nowcasting from Radar Images. In.

Andrews, D. G. (1987). The influence of atmospheric waves on the general circulation of the middle atmosphere. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 323(1575), 693-705. <https://doi.org/10.1098/rsta.1987.0115>

Arakawa, A. (2004). The Cumulus Parameterization Problem: Past, Present, and Future. *Journal of Climate*, 17(13), 2493-2525. [https://doi.org/10.1175/1520-0442\(2004\)017<2493:Ratcpp>2.0.Co;2](https://doi.org/10.1175/1520-0442(2004)017<2493:Ratcpp>2.0.Co;2)

Arakawa, A., & Schubert, W. H. (1974). Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I. In *Journal of the Atmospheric Sciences* (Vol. 31, pp. 674-701).

Baldwin, M. P., Ayarzagüena, B., Birner, T., Butchart, N., Butler, A. H., Charlton-Perez, A. J., Domeisen, D. I. V., Garfinkel, C. I., Garny, H., Gerber, E. P., Hegglin, M. I., Langematz, U., & Pedatella, N. M. (2021). Sudden Stratospheric Warmings. *Reviews of Geophysics*, 59(1). <https://doi.org/10.1029/2020rg000708>

Baldwin, M. P., Stephenson, D. B., Thompson, D. W., Dunkerton, T. J., Charlton, A. J., & O'Neill, A. (2003). Stratospheric memory and skill of extended-range weather forecasts. *Science*, 301(5633), 636-640.

Beucler, T., Pritchard, M., Rasp, S., Ott, J., Baldi, P., & Gentine, P. (2021). Enforcing Analytic Constraints in Neural Networks Emulating Physical Systems. *Physical Review Letters*, 126(9), 098302. <https://doi.org/10.1103/PhysRevLett.126.098302>

Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2023). Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970), 533-538. <https://doi.org/10.1038/s41586-023-06185-3>

Böing, S. J. (2016). An object-based model for convective cold pool dynamics. In *Mathematics of Climate and Weather Forecasting* (Vol. 2, pp. 43-60).

Brüning, S., Niebler, S., & Tost, H. (2024). Artificial intelligence (AI)-derived 3D cloud tomography from geostationary 2D satellite data. *Atmos. Meas. Tech.*, 17(3), 961-978. <https://doi.org/10.5194/amt-17-961-2024>

Charlton, A. J., & Polvani, L. M. (2007). A New Look at Stratospheric Sudden Warmings. Part I: Climatology and Modeling Benchmarks. *Journal of Climate*, 20(3), 449-469. <https://doi.org/10.1175/jcli3996.1>

Chattopadhyay, A., Hassanzadeh, P., & Pasha, S. (2020). Predicting clustered weather

patterns: A test case for applications of convolutional neural networks to spatio-temporal climate data. In *Scientific Reports 2020 10:1* (Vol. 10, pp. 1-13): Nature Publishing Group.

Chen, P.-Y., & Wu, C.-M. (2025). Identifying Cold Pool Scales over Complex Topography Using TaiwanVVM Simulations. *Journal of the Meteorological Society of Japan. Ser. II*. <https://doi.org/10.2151/jmsj.2025-023>

Chen, Y.-C., Liang, Y.-C., Wu, C.-M., Huang, J.-D., Lee, S. H., Wang, Y., & Zeng, Y.-J. (2024). Exploiting a variational auto-encoder to represent the evolution of sudden stratospheric warmings. *Environmental Research: Climate*, 3(2). <https://doi.org/10.1088/2752-5295/ad3a0d>

Chen, Y. C., & Wu, C. M. (2025). Capturing Subgrid Cold Pool Dynamics With U-Net: Insights From Large-Eddy Simulation for Storm-Resolving Modeling. *Atmospheric Science Letters*, 26(7). <https://doi.org/10.1002/asl.1309>

Chen, Y. C., Wu, C. M., & Chen, W. T. (2023). A deep learning framework for analyzing cloud characteristics of aggregated convection using cloud-resolving model simulations. *Atmospheric Science Letters*. <https://doi.org/10.1002/asl.1150>

Davis, N., Richter, J., Glanville, A., Edwards, J., & LaJoie, E. (2022). Limited surface impacts of the January 2021 sudden stratospheric warming. *Nature communications*, 13(1), 1136.

Domeisen, D. I., Butler, A. H., Charlton-Perez, A. J., Ayarzagüena, B., Baldwin, M. P., Dunn-Sigouin, E., Furtado, J. C., Garfinkel, C. I., Hitchcock, P., & Karpechko, A. Y. (2020). The role of the stratosphere in subseasonal to seasonal prediction: 2. Predictability arising from stratosphere-troposphere coupling. *Journal of Geophysical Research: Atmospheres*, 125(2), e2019JD030923.

Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5), 1937-1958.

Feng, Z., Hagos, S., Rowe, A. K., Burleyson, C. D., Martini, M. N., & De Szoke, S. P. (2015). Mechanisms of convective cloud organization by cold pools over tropical warm ocean during the AMIE/DYNAMO field campaign. In *Journal of Advances in Modeling Earth Systems* (Vol. 7, pp. 357-381): Blackwell Publishing Ltd.

Gelaro, R., McCarty, W., Suárez, M. J., Todling, R., Molod, A., Takacs, L., Randles, C. A., Darmenov, A., Bosilovich, M. G., & Reichle, R. (2017). The modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *Journal of Climate*, 30(14), 5419-5454.

Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., & Yacalis, G. (2018). Could Machine Learning Break the Convection Parameterization Deadlock? In *Geophysical*

Research Letters (Vol. 45, pp. 5742-5751): Blackwell Publishing Ltd.

Gettelman, A., Mills, M., Kinnison, D., Garcia, R., Smith, A., Marsh, D., Tilmes, S., Vitt, F., Bardeen, C., & McInerny, J. (2019). The whole atmosphere community climate model version 6 (WACCM6). *Journal of Geophysical Research: Atmospheres*, 124(23), 12380-12403.

Haarsma, R. J., Roberts, M. J., Vidale, P. L., Senior, C. A., Bellucci, A., Bao, Q., Chang, P., Corti, S., Fučkar, N. S., & Guemas, V. (2016). High resolution model intercomparison project (HighResMIP v1. 0) for CMIP6. *Geoscientific Model Development*, 9(11), 4185-4208.

Haerter, J. O. (2019). Convective Self-Aggregation As a Cold Pool-Driven Critical Phenomenon. In *Geophysical Research Letters* (Vol. 46, pp. 4017-4028).

Haerter, J. O., Böing, S. J., Henneberg, O., & Nissen, S. B. (2019). Circling in on Convective Organization. *Geophysical Research Letters*, 46(12), 7024-7034. <https://doi.org/10.1029/2019gl082092>

Ham, Y.-G., Kim, J.-H., Min, S.-K., Kim, D., Li, T., Timmermann, A., & Stuecker, M. F. (2019). Deep learning for multi-year ENSO forecasts. *Nature*, 573(7775), 568-572. <https://doi.org/10.1038/s41586-019-1559-7>

Hamada, A., Murayama, Y., & Takayabu, Y. N. (2014). Regional Characteristics of Extreme Rainfall Extracted from TRMM PR Measurements. *Journal of Climate*, 27(21), 8151-8169. <https://doi.org/10.1175/jcli-d-14-00107.1>

Higa, M., Tanahara, S., Adachi, Y., Ishiki, N., Nakama, S., Yamada, H., Ito, K., Kitamoto, A., & Miyata, R. (2021). Domain knowledge integration into deep learning for typhoon intensity classification. In *Scientific Reports* (Vol. 11, pp. 1-10): Nature Publishing Group UK.

Hoeller, J., Fiévet, R., & Haerter, J. O. (2022). U-Net Segmentation for the Detection of Convective Cold Pools From Cloud and Rainfall Fields. <https://doi.org/10.1002/essoar.10512898.1>

Hohenegger, C., & Schlemmer, L. (2014). The Formation of Wider and Deeper Clouds as a Result of Cold-Pool Dynamics. *Journal of the Atmospheric Sciences*, 71(8), 2842-2858. <https://doi.org/10.1175/jas-d-13-0170.1>

Iacono, M. J., Delamere, J. S., Mlawer, E. J., Shephard, M. W., Clough, S. A., & Collins, W. D. (2008). Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models. *Journal of Geophysical Research: Atmospheres*, 113(D13). <https://doi.org/10.1029/2008jd009944>

Jung, J. H., & Arakawa, A. (2008). A three-dimensional anelastic model based on the vorticity equation. In *Monthly weather review* (Vol. 136, pp. 276-294).

Kao, F. S., Kuo, Y. H., & Wu, C. M. (2025). Diagnosing Nonlocal Vertical Acceleration in Moist Convection Using a Large-Eddy Simulation. *Journal of Advances in*

Modeling Earth Systems, 17(3). <https://doi.org/10.1029/2024ms004636>

Khairoutdinov, M., & Randall, D. (2006). High-resolution simulation of shallow-to-deep convection transition over land. In *Journal of the Atmospheric Sciences* (Vol. 63, pp. 3421-3436).

Kidston, J., Scaife, A. A., Hardiman, S. C., Mitchell, D. M., Butchart, N., Baldwin, M. P., & Gray, L. J. (2015). Stratospheric influence on tropospheric jet streams, storm tracks and surface weather. *Nature Geoscience*, 8(6), 433-440.

Kim, K.-H., Shim, P.-S., & Shin, S. (2019). An Alternative Bilinear Interpolation Method Between Spherical Grids. *Atmosphere*, 10(3). <https://doi.org/10.3390/atmos10030123>

Kingma, D. P., & Welling, M. (2013). Auto-Encoding Variational Bayes. arXiv:1312.6114. Retrieved December 01, 2013, from <https://ui.adsabs.harvard.edu/abs/2013arXiv1312.6114K>

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., & Battaglia, P. (2023). Learning skillful medium-range global weather forecasting. *Science*, 382(6677), 1416-1421. <https://doi.org/doi:10.1126/science.adf2336>

Lehtonen, I., & Karpechko, A. Y. (2016). Observed and modeled tropospheric cold anomalies associated with sudden stratospheric warmings. *Journal of Geophysical Research: Atmospheres*, 121(4), 1591-1610. <https://doi.org/https://doi.org/10.1002/2015JD023860>

Liang, Y.-C., Kwon, Y.-O., Frankignoul, C., Danabasoglu, G., Yeager, S., Cherchi, A., Gao, Y., Gastineau, G., Ghosh, R., Matei, D., Mecking, J. V., Peano, D., Suo, L., & Tian, T. (2020). Quantification of the Arctic Sea Ice-Driven Atmospheric Circulation Variability in Coordinated Large Ensemble Simulations. *Geophysical Research Letters*, 47(1), e2019GL085397. <https://doi.org/https://doi.org/10.1029/2019GL085397>

Liang, Z., Rao, J., Guo, D., & Lu, Q. (2022). Simulation and projection of the sudden stratospheric warming events in different scenarios by CESM2-WACCM. *Climate Dynamics*, 59(11-12), 3741-3761.

Limpasuvan, V., Thompson, D. W., & Hartmann, D. L. (2004). The life cycle of the Northern Hemisphere sudden stratospheric warmings. *Journal of Climate*, 17(13), 2584-2596.

Matthewman, N. J., Esler, J. G., Charlton-Perez, A. J., & Polvani, L. M. (2009). A new look at stratospheric sudden warmings. Part III: Polar vortex evolution and vertical structure. *Journal of Climate*, 22(6), 1566-1585.

Mitchell, D. M., Charlton-Perez, A. J., & Gray, L. J. (2011). Characterizing the variability

and extremes of the stratospheric polar vortices using 2D moment analysis. *Journal of the Atmospheric Sciences*, 68(6), 1194-1213.

Morrison, H., & Milbrandt, J. A. (2015). Parameterization of Cloud Microphysics Based on the Prediction of Bulk Ice Particle Properties. Part I: Scheme Description and Idealized Tests. *Journal of the Atmospheric Sciences*, 72(1), 287-311. <https://doi.org/10.1175/jas-d-14-0065.1>

Moseley, C., Hohenegger, C., Berg, P., & Haerter, J. O. (2016). Intensification of convective extremes driven by cloud–cloud interaction. *Nature Geoscience*, 9(10), 748-752. <https://doi.org/10.1038/ngeo2789>

North, G. R., Bell, T. L., Cahalan, R. F., & Moeng, F. J. (1982). Sampling errors in the estimation of empirical orthogonal functions. *Monthly weather review*, 110(7), 699-706.

Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684-9689. <https://doi.org/doi:10.1073/pnas.1810286115>

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?" Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. *ArXiv*, *abs/1505.04597*.

Roscher, R., Bohn, B., Duarte, M. F., & Garcke, J. (2020). Explainable Machine Learning for Scientific Insights and Discoveries. *IEEE Access*, 8, 42200-42216. <https://doi.org/10.1109/access.2020.2976199>

Rotunno, R., Klemp, J. B., & Weisman, M. L. (1988). A Theory for Strong, Long-Lived Squall Lines. *Journal of the Atmospheric Sciences*, 45(3), 463-485. [https://doi.org/10.1175/1520-0469\(1988\)045<0463:Atfsll>2.0.Co;2](https://doi.org/10.1175/1520-0469(1988)045<0463:Atfsll>2.0.Co;2)

Samek, W., Wiegand, T., & Müller, K.-R. (2017). Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models. In.

Seviour, W. J. M., Mitchell, D. M., & Gray, L. J. (2013). A practical method to identify displaced and split stratospheric polar vortex events. *Geophysical Research Letters*, 40(19), 5268-5273. <https://doi.org/10.1002/grl.50927>

Shutts, G. J., & Gray, M. E. B. (2006). A numerical modelling study of the geostrophic adjustment process following deep convection. *Quarterly Journal of the Royal Meteorological Society*, 120(519), 1145-1178. <https://doi.org/10.1002/qj.49712051903>

Sigmond, M., Scinocca, J., Kharin, V., & Shepherd, T. (2013). Enhanced seasonal forecast skill following stratospheric sudden warmings. *Nature Geoscience*, 6(2), 98-102.

Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale

image recognition. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings* (pp. 1-14).

Stevens, B., Satoh, M., Auger, L., Biercamp, J., Bretherton, C. S., Chen, X., Düben, P., Judt, F., Khairoutdinov, M., Klocke, D., Kodama, C., Kornblueh, L., Lin, S.-J., Neumann, P., Putman, W. M., Röber, N., Shibuya, R., Vanniere, B., Vidale, P. L.,...Zhou, L. (2019). DYAMOND: the DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains. *Progress in Earth and Planetary Science*, 6(1), 61. <https://doi.org/10.1186/s40645-019-0304-z>

Tompkins, A. M. (2001). Organization of Tropical Convection in Low Vertical Wind Shears: The Role of Cold Pools. *Journal of the Atmospheric Sciences*, 58(13), 1650-1672. [https://doi.org/10.1175/1520-0469\(2001\)058<1650:Ootcil>2.0.Co;2](https://doi.org/10.1175/1520-0469(2001)058<1650:Ootcil>2.0.Co;2)

Torri, G., & Kuang, Z. (2019). On Cold Pool Collisions in Tropical Boundary Layers. *Geophysical Research Letters*, 46(1), 399-407. <https://doi.org/10.1029/2018gl080501>

Tripathi, O. P., Baldwin, M., Charlton-Perez, A., Charron, M., Eckermann, S. D., Gerber, E., Harrison, R. G., Jackson, D. R., Kim, B. M., & Kuroda, Y. (2015). The predictability of the extratropical stratosphere on monthly time-scales and its impact on the skill of tropospheric forecasts. *Quarterly Journal of the Royal Meteorological Society*, 141(689), 987-1003.

Tripathi, O. P., Charlton-Perez, A., Sigmond, M., & Vitart, F. (2015). Enhanced long-range forecast skill in boreal winter following stratospheric strong vortex conditions. *Environmental Research Letters*, 10(10), 104007.

Tsai, J.-Y., & Wu, C.-M. (2016). Critical transitions of stratocumulus dynamical systems due to perturbation in free-atmosphere moisture. *Dynamics of Atmospheres and Oceans*, 76, 1-13. <https://doi.org/https://doi.org/10.1016/j.dynatmoce.2016.08.002>

Tsai, W. M., & Wu, C. M. (2017). The environment of aggregated deep convection. In *Journal of Advances in Modeling Earth Systems* (Vol. 9, pp. 2061-2078).

Tsou, S.-W., Su, C.-Y., & Wu, C.-M. (2019). Learning the Representations of Moist Convection with Convolutional Neural Networks. In (pp. 1-5).

Wang, Y.-H., Chen, W.-T., & Wu, C.-M. (2024). Novel perspectives on multiple-peak diurnal convection over a tropical mountainous island from idealized large-eddy simulations. *npj Climate and Atmospheric Science*, 7(1). <https://doi.org/10.1038/s41612-024-00884-y>

Weyn, J. A., Durran, D. R., & Caruana, R. (2020). Improving Data-Driven Global Weather Prediction Using Deep Convolutional Neural Networks on a Cubed Sphere. *Journal of Advances in Modeling Earth Systems*, 12(9). <https://doi.org/10.1029/2020ms002109>

Wu, C.-M., & Arakawa, A. (2011). Inclusion of Surface Topography into the Vector Vorticity Equation Model (VVM). *Journal of Advances in Modeling Earth Systems*, 3(2), n/a-n/a. <https://doi.org/10.1029/2011ms000061>

Wu, C. M., Lin, H. C., Cheng, F. Y., & Chien, M. H. (2019). Implementation of the Land Surface Processes into a Vector Vorticity Equation Model (VVM) to Study its Impact on Afternoon Thunderstorms over Complex Topography in Taiwan. In *Asia-Pacific Journal of Atmospheric Sciences* (Vol. 55, pp. 701-717): Korean Meteorological Society.

Xu, K. M., Hu, Y., & Wong, T. (2019). Convective Aggregation and Indices Examined from CERES Cloud Object Data. In *Journal of Geophysical Research: Atmospheres* (Vol. 124, pp. 13604-13624).

Yuval, J., & O'Gorman, P. A. (2020). Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions. *Nat Commun*, 11(1), 3295. <https://doi.org/10.1038/s41467-020-17142-3>