

Generative Modeling of Atmospheric Convection

Griffin Mooers*
gmooers96@gmail.com
University of California Irvine
Irvine, California, USA

Jens Tuyls
University of California Irvine
Irvine, USA
jtuyls@uci.edu

Stephan Mandt
University of California Irvine
Irvine, USA
mandt@uci.edu

Mike Pritchard
University of California Irvine
Irvine, USA
mspritch@uci.edu

Tom Beucler
University of California Irvine
Irvine, USA
tbeucler@uci.edu

ABSTRACT

While cloud-resolving models can explicitly simulate the details of small-scale storm formation and morphology, these details are often ignored by climate models for lack of computational resources. Here, we explore the potential of generative modeling to cheaply recreate small-scale storms by designing and implementing a Variational Autoencoder (VAE) that performs structural replication, dimensionality reduction, and clustering of high-resolution vertical velocity fields. Trained on $\sim 6 \cdot 10^6$ samples spanning the globe, the VAE successfully reconstructs the spatial structure of convection, performs unsupervised clustering of convective organization regimes, and identifies anomalous storm activity, confirming the potential of generative modeling to power stochastic parameterizations of convection in climate models.

KEYWORDS

variational autoencoders, climate modeling, vertical velocity, sub-grid parameterization,

ACM Reference Format:

Griffin Mooers, Jens Tuyls, Stephan Mandt, Mike Pritchard, and Tom Beucler. 2020. Generative Modeling of Atmospheric Convection. In *10th International Conference on Climate Informatics (CI2020)*, September 22–25, 2020, virtual, United Kingdom. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3429309.3429324>

1 INTRODUCTION

Boxed in by computational limits, many of the details of our atmosphere remain too minute to explicitly resolve in climate models [14, 27, 29]. Key physics driving convection and cloud formation occur on the scale of meters to a few kilometers, while typical modern climate models have a resolution of $100 - 200\text{km}^2$ horizontally - meaning important sub-grid processes are parameterized. Computational capabilities are advancing, and climate models are increasingly common, in particular those with three-dimensional explicit resolution of clouds systems. However, the capability to

run these models for the ~ 100 -year timescales needed is often impractical [13, 15, 21] and the information content they generate about the details of cloud and storm organization are frequently overwhelming to analyze at its native scale. This has left significant gaps in knowledge about many of the details of cloud-climate feedbacks and the relationship between storm organization and its thermodynamic environment [21, 27]. However, deep learning, and in particular generative models, may provide a path to a better understanding of these phenomena and their role driving the weather and climate of our world.

The application of machine learning in the physical sciences has increased exponentially in recent years but with important avenues still largely unexplored. In climate modeling, deep neural networks have been re-purposed to emulate the large-scale consequences of storm-level heating and moistening over the atmospheric column to replicate mean climate and expected precipitation patterns and extremes [8, 21, 24, 25, 28]. However, much of this work has been confined to deterministic neural networks that ignore the interesting stochastic details of eddy and storm organization. The recent application of Generative Adversarial Networks (GANs, [10]) to the Lorenz '96 Model suggests a potential, under-explored role for generative models in atmospheric sciences - particularly towards stochastic parameterizations [5, 7]. There have also been initial successes using various types of GAN architectures to generate plausible Rayleigh-Bernard convection. In particular, adding informed physical constraints to GAN loss functions seem to improve the generation of these non-linear fluid flow systems [30, 33, 35, 36]. While promising, such techniques have thus far been restricted to idealized turbulent flows of reduced dimension and complexity; there is ample room to explore generative modeling methods for representing convective details amidst settings of realistic geographic complexity. Meanwhile, generative modeling besides GANs have not been as thoroughly considered for turbulent flow emulation and could potentially power climate models down the line.

VAEs may prove more appropriate than GANs for these climate applications given their design containing both a generative and representational model, their often superior log-likelihoods and reconstruction simulations, and practical advantages including stabler training results, easier performance benchmarking, and more interpretable latent manifold representations [12, 22, 34]. Modified VAEs can reconstruct plausible two-dimensional laminar flow with computational efficiency beyond what is common when numerically solving linear differential equations [6]. There has been preliminary



This work is licensed under a Creative Commons Attribution International 4.0 License.
CI2020, September 22–25, 2020, virtual, United Kingdom
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8848-1/20/09.
<https://doi.org/10.1145/3429309.3429324>

Layer	Filters	Kernel	Stride	Activation
2D Conv	64	3x3	2	relu
2D Conv	128	3x3	2	relu
2D Conv	512	3x3	2	relu
2D Conv (μ)	64	3x3	2	relu
2D Conv (σ)	64	3x3	2	relu

Table 1: Our Encoder architecture. Conv refers to a convolutional hidden layer. The first hidden Conv layer receives an input vector of 32x128 (30x128 expanded by padding) representing a vertical velocity snapshot.

work using VAEs for the clustering of atmospheric dynamics – a gain again relying on simplified Lorenz ’96 model data as well as potential vorticity fields and geopotential heights [31, 32]. This application of representation learning across a variety of simplified simulations suggests VAEs offer great potential as both an engineering tool to help escape computational limits on the generative side and may provide the ability to learn and extract hidden organizational details in atmospheric dynamics on the representation side. However, to the best of our knowledge, this is the first study to use a VAE for representational learning on the details of convective organization and associated gravity wave radiation¹ as revealed by spatial snapshots of vertical velocity – an inherently chaotic and bimodal variable [9] – across a dataset large enough to nonetheless encompass the spatiotemporal diversity of turbulence regimes in the atmosphere. As far as we know, this is also the first study to constrain a VAE’s output statistics by adding a covariance constraint term to its loss function to improve representation and capture variance details at small spatial scales in the turbulent atmospheric boundary layer, which can be considered one of the most difficult locations for climate models. Our results demonstrate the power of VAEs to accurately reconstruct high-resolution climate data, as well as their ability to leverage dimensionality reduction for high level feature learning and anomaly detection. This casts VAEs as promising tools for both dynamical analysis and stochastic parameterization of fine-scale atmospheric processes from cloud-resolving data.

2 METHODS

In this Section, we discuss the architecture of the three machine-learning models used here, the design of our covariance constrained VAE loss function, and the generation and preprocessing of the atmospheric simulation data.

2.1 Architecture

Our VAE takes vertical velocity fields formatted as (30x128) 2D images. We adopt a fully convolutional design² to preserve local information, which is essential in atmospheric convection modeling

¹Here we are referring to internal gravity waves, which are horizontally-propagating disturbances in the atmosphere generated by density perturbations, e.g. from deep convection, frontogenesis, or topography.

²Earlier experiments used architecture similar to models used for CIFAR-10 data [19] with fully connected dense layers separating the encoder and the decoder from the latent space, but led to discouraging reconstructions plagued by posterior collapse and an inability to represent the spatial patterns of convection.

Layer	Filters	Kernel	Stride	Activation
2D Conv-T	1024	3x3	2	relu
2D Conv-T	256	3x3	2	relu
2D Conv-T	64	3x3	2	relu
2D Conv (μ)	1	3x3	2	sigmoid
2D Conv (σ)	1	3x3	2	linear

Table 2: Our Decoder architecture. Conv-T refers to a transposed convolutional hidden layer.

(Tables 1 and 2). We obtain meaningful reconstruction performance by ensuring that the information bottleneck in the VAE is not too severe, i.e. that the latent space is still wide enough to preserve enough fine features of the vertical velocity fields (in our case of dimension 1024), and by implementing annealing techniques outlined in [1, 11]. Here, we analyze two successful VAEs: One with a traditional negative ELBO in the loss, and one with an additional covariance constraint in the loss. As a baseline, we also implemented a regular autoencoder of the same design as above, with two key differences: All activations were replaced with the identity function and our covariance constrained loss was replaced with the mean-squared error. We refer to this model as the “linear” model, and use it to better quantify the added value of VAEs for modeling atmospheric convection.

2.2 VAE Loss Implementation

The total loss is the sum of two terms: the negative of the Evidence Lower Bound (ELBO), commonly used as the total VAE loss, and a covariance constraint loss term [6, 20, 30] on the covariance matrix that we weigh by $\lambda \in \mathbb{R}^+$:

$$\text{Loss} \stackrel{\text{def}}{=} -\text{ELBO} + \lambda \times \text{CC}, \quad (1)$$

where CC is a “covariance constraining” term using the Frobenius norm $\|\cdot\|$ to measure the distance between the covariance, Σ , of the likelihood $p_\theta(\mathbf{x}|\mathbf{z})$ and the covariance, Σ , of the true data distribution $p(\mathbf{x})$. θ refers to model parameters and $\mathbf{x}$ refers to observed vertical velocity fields:

$$\text{CC} = \|\Sigma(p_\theta(\mathbf{x}|\mathbf{z})) - \Sigma(p(\mathbf{x}))\|. \quad (2)$$

Unconstrained VAEs ($\lambda = 0$), henceforth referred to as “VAE” for short, maximize the ELBO, defined as the sum of the log-likelihood $p_\theta(\mathbf{x}|\mathbf{z})$, and the Kullback-Leibler (KL) Divergence between $p(\mathbf{z})$ and $q_\phi(\mathbf{z}|\mathbf{x})$:

$$\begin{aligned} \text{ELBO}(\mathbf{x}; \theta, \phi, \mathbf{z}) = & \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] \\ & - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})), \end{aligned} \quad (3)$$

where ϕ are our variational parameters which are learned jointly with the model parameters, θ . $p(\mathbf{z})$ refers to the prior and $q_\phi(\mathbf{z}|\mathbf{x})$ refers to the estimated posterior. We denote hidden variables as $\mathbf{z}$. Minimizing the KL loss term regularizes the variational parameters in the model and makes the VAE posterior more similar to the VAE prior. Maximizing the log-likelihood enables the VAE to produce realistic vertical velocity fields where the output will be more closely aligned with the latent variable of the model. Following [18], we

assume that the prior over the parameters and the hidden variables are both centered isotropic Gaussian and calculate ELBO using equation (24) of [18].

To control the rate-distortion trade-off [1], we implement linear annealing to the KL loss term following [2], where the KL term is multiplied by an annealing factor linearly scaled from 0 to 1 over the course of training. In our VAE, linear annealing results in significantly lower KL losses and more interpretable latent spaces.

Finally, to generate vertical velocity fields with realistic spatial variability, we additionally implement covariance-constrained VAEs. Following Equation 2, the covariance constraint is defined as the Frobenius norm of the covariance matrix error, which we estimate over each batch during optimization. We choose a pre-factor $\lambda = 10^6$ so that the magnitude of the covariance constraint matches that of the reconstruction loss, resulting in a covariance-constrained VAE “CC-VAE” that generates more faithful covariance matrices.

2.3 Data & Preprocessing

2.3.1 Cloud-Resolving Data. To train and test our VAE, we rely on snapshots of vertical motions with explicitly-resolved moist convection and gravity wave radiation obtained from $\sim 15k$ instances of a Cloud-Resolving Model (CRM) [16, 17] embedded within a host Global Climate Model (GCM). The CRMs operates at a 20s native timestep data and we extract state snapshots from it every 15 minutes, the frequency with which its horizontal average state is permitted to interact with its host GCM. We perform a 100-day multi-scale climate simulation to generate data showing details of atmospheric convection within a tropical belt from 20N to 20S latitudes. Specifically, at each $1.9^\circ \times 2.5^\circ$ horizontal grid cell of the Super-Parameterized Community Atmosphere Model (SPCAM5), we embed a 128-column System for Atmospheric Modeling (SAM) micro model with kilometer scale horizontal resolution; both the host and embedded models use 30 vertical levels. This entire dataset comes to a size of 1.3 Tb. For our purposes, there is 30 level by 128 CRM-column “snapshot” or “image” of a convective-scale vertical velocity field at each latitude-longitude grid cell that we feed into the encoder of our neural network. We train our VAEs on sub-samples of this data staged on UC Irvine’s GreenPlanet Super-computing node and our machine learning simulations are powered by two NVIDIA Tesla V100 and one NVIDIA Tesla T4 GPUs.

2.3.2 Preprocessing. To reduce data volume for efficient training and to ensure our VAE is exposed to a plethora of convective motion, we selectively sample from the initial 1.3Tb SAM dataset. We restrict our initial data volume to the 144 latitude/longitude coordinates with a detectable diurnal cycle of precipitation where amplitude of daily precipitation is greater than two times its standard deviation within the larger-scale host model. This precipitation filtering ensures samples of strong convection get placed into the training dataset, as a persistent diurnal cycle of precipitation often indicates deep convection and the presence of mesoscale convective systems [3]. Within these selected grid cells, the vertical velocity values range from $37.3m\ s^{-1}$ to $-17.4m\ s^{-1}$ and are then scaled from 0-1 by subtracting the minimum and dividing by the range.

We shuffle data in the spatial and temporal dimensions prior to training. We use An 80%/20% training/test split for all models. To

ensure a balanced dataset of different convective types, we apply K-means clustering with two centroids to group data with active and inactive vertical velocity fields. We then sample equally from both clusters without replacement to design a balanced dataset for the VAE. This new 4.3Gb dataset has a 111206/27802 training/test split. Since the horizontal domain is doubly-periodic, two vertical velocity updrafts of equal magnitudes and size located at different horizontal locations are physically identical. To prevent the VAE from treating them as different at the expense of reconstruction magnitude and variance, we preprocess all samples so that the center of the vertical velocity field is the location of strongest convection present in the sample. We define the “strongest convection” as the largest absolute value of spatially-averaged vertical velocity, from 400hPa to 600hPa in the vertical and using a moving average of 10km horizontally.

2.4 Quantifying Reconstruction Performance

We quantify the reconstructions of our final VAE and CC VAE as well as our linear baseline using the following metrics:

2.4.1 Hellinger Distance. We calculate the Hellinger distance H between the discrete distributions to gauge similarity [23]:

$$H(p, q) = \sqrt{\sum_{i=1}^k \frac{(\sqrt{p_i} - \sqrt{q_i})^2}{2}} \quad (4)$$

where p is the distribution of the original vertical velocity fields and q is the distribution of the corresponding reconstruction.

2.4.2 Mean Squared Error (MSE). To provide an overall skill of the reconstruction, the MSE is calculated between each original sample and its corresponding reconstruction.

2.4.3 Spectral Analysis. To better understand the skill of the VAE reconstruction from a spatial perspective, we perform one-dimensional spectral analysis on each sample and reconstruction at all 30 levels in the vertical dimension. We examine four vertical levels commonly used in meteorology: 850hPa (top of the boundary layer), 700hPa (lower troposphere), 500hPa (mid-troposphere), and 250hPa (upper-troposphere) to see how our VAEs capture the spatially-resolved vertical velocity variance throughout the atmosphere. We calculate the power spectral density Φ_k using:

$$\Phi_k \stackrel{\text{def}}{=} \frac{\Delta n}{N} \left| \sum_{j=0}^{N-1} y_j e^{-\frac{ijk}{NT}} \right|^2 \quad (5)$$

where N is the length of the x dimension, y_j is the sample or reconstruction, T is $1/\text{length}$, i is the imaginary unit and k is the vertical level of interest in hPa (850, 700, 500, or 250) [4].

3 RESULT & DISCUSSION

Our VAE trained on cloud-resolving climate data produces accurate vertical velocity field reconstructions. When we provide the high resolution training dataset and appropriate convolutional architecture, our VAE learns remarkably accurate representations of any type of convection found within the test dataset. Our VAE captures the magnitude, proper height, and structure across deep convective regimes, shallow convective regimes, and non-convecting regimes (Figure 4). When the “Covariance Constraining” term is added to




Figure 1: Visualization of the latent space originally in dimension 1024, but reduced to dimension 2 by Principle Component Analysis (PCA) [26]. The standard deviations of different types of convection the VAE learns to cluster are embedded near corresponding clusters. This suggests the VAE learns an interpretable clustering of the data, with means and variances both contributing to the results.




Figure 2: Spectral Analysis at 4 different levels of the atmosphere comparing the test data to our best VAE and CC VAE as well as a linear model. At small spatial scales we see the importance of the Covariance Constraint to capture the variance native to convection (orange vs. red).

create a physically informed loss, the CC VAE performance improves enough to match a linear baseline (Table 3). But unlike many other image recognition tasks generative models perform, reconstructing the mean of the convection is necessary but not sufficient – we must capture the variance and correlation in the vertical velocity fields. The CC VAE reconstructs variance better than a traditional convolutional VAE and at least on par with the linear baseline (Table 3, Figure 2). Our CC VAE is the most versatile of our models with an accurate reconstruction performance overall at different levels of the atmospheric column and different convective spatial scales based on the power spectra of the three models (Figure 2). This precision across both small and large spatial scales revealing our CC VAEs ability to emulate both the overall large

pattern of convective plumes and the details within convective composition. Our CC VAEs results replicate disparate structures of convection in both areas of high stochasticity near the atmospheric boundary layer, characteristic of shallow convection, as well as in the upper troposphere, where deep convective regimes dominate. At this stage CC VAEs match the performance of our linear baseline but do not exceed it.

However, unlike the linear baseline, our VAE and CC VAE discover the details of convective organization by representation learning via dimensionality reduction and feature extraction. A 2D, deterministic PCA projection of our CC VAE latent space clusters and separates different convective types (Figure 1). In particular, the distinction between deep and shallow convective regimes and




Figure 3: Convection Type Predictions The diurnal composite from a ten day average at four unique times of day are shown above. The VAE predicts the type of convection occurring in tropical locations over the course of a typical Boreal Winter Diurnal Cycle. Semantic similarities of the VAE latent space are reflected in the global geospatial weather patterns.




Figure 4: Reconstructions The trained VAE reconstructions closely resemble those from the test dataset and accurately predict the location, magnitude and spatial structure of convective plumes.




Figure 5: 2D PCA Temporal Projection All spatial locations comprising the Amazon Rainforest are averaged together from November to February to get a single composite diurnal cycle that is fed through our trained VAE. The colors above correlate to time of day (Local Solar Time). The results show a clear separation in representation on the latent space of the timing of deepest convection and maximum precipitation (mid afternoon) from when shallow convection and calmer conditions dominate (early morning).




Figure 6: Anomaly Detection We use the ELBO in the VAE Loss function to identify the most anomalous vertical velocity fields. We show the 9th most anomalous field because it exhibits multiple deep convective plumes.

Model	MSE	Hellinger Distance	Frobenius Norm
Linear	4.2e-6	2.0e-3	8.0e-3
VAE	1.1e-5	3.1e-4	3.2e-4
CC VAE	4.5e-6	2.0e-3	8.0e-6

Table 3: Quantitative Reconstruction Metrics. We compute the MSE and Hellinger Distance between true and predicted reconstructions. This shows the baseline is equally good at predicting the mean reconstruction. We also compute the Frobenius Norm of the error in the covariance matrices of the true data and the reconstructions. Both VAEs capture more of the covariance structure of the data than the linear baseline.

non-convective regimes is encouraging (Figure 1, please visit this link for a complete animation of the 2D Projection of the latent space). The physical knowledge represented in our CC VAEs latent space stands alone from other forms of dimensionality reduction (PCA and t-SNE on the preprocessed data) where there is no evidence of distinction based on convective type. Furthermore, CC VAE predictions of convective type based solely on latent space location map back to a physically sensible pattern over the tropics with deep convection concentrated on land over the Amazon and African Rainforests as well as over the Pacific Warmpool (Figure 3). These predictions from latent space location not only map convection type in a spatially coherent pattern, but also capture the change in convection type with the diurnal cycle over moist, tropical continents (Figure 3, please visit this link for a complete animation of the tropical diurnal cycle). When we exclusively restrict the test dataset to an Amazon Diurnal Composite, the known coherent transitions from shallow to deep convection that occur over tropical rain-forest in response to solar heating of the diurnal cycle correspond to monotonic trajectories in the latent space projection, verified using both t-SNE and PCA (Figure 5). Further tests are required on more complex convective transitions to understand

the extent of the physical meaning of the CC VAE latent space, but these initial positive results suggest great potential for physically constrained VAEs as a tool in atmospheric dynamics to uncover information about convective transitions, storm morphology and propagation.

We also evaluate ELBO (Equation 3) for each sample of our test data to find unusual storm development and activity in the dense CRM data.

ELBO allows us to determine the degree to which a vertical velocity field, drawn from our models latent variables is an aberration in the data. Our VAEs inherent ability to detect anomalies in the vertical velocity data proves to be an elegant way to identify deep convection in a more thorough manner than traditional vertical velocity thresholding. An example of one such anomaly we identify is Figure 6 – in this case an instance of two moderate storms developing in one CRM array. This phenomena would be less straightforward to locate through conventional methods, particularly given the size and density of data involved. Our VAEs attribute of anomaly detection learns characteristics of the data instead of naively thresholding based on priors experiences that may not reflect the composition of the dataset. This feature provides the potential to help identify interesting and unexpected weather phenomena from noise – artifacts that might otherwise never be studied in overwhelmingly large and rich datasets.

4 CONCLUSION

We develop a VAE to reconstruct immaculate convection images from a high-resolution, cloud-resolving dataset. Our VAE, particularly once a statistically constrained loss function is added, captures the variance and magnitude of distinct convective regimes. The latent space of the VAE proves to be a potent tool for making physically sensible predictions of convection type that accurately reflect the tropical atmosphere and capture the effects of solar heating through the diurnal cycle. The unique VAE loss function allows us to use ELBO to find anomalous storm development in a dense, high resolution dataset that traditional methods might miss. But

there is much work to be done before a VAE could be implemented to power stochastic parameterizations for a climate model, likely requiring to condition the VAE on large-scale thermodynamics via expansion of the input vector. If successful, the ability to quickly and efficiently generate synthetic, detailed vertical velocity fields to help run climate models would be a valuable resource for the atmospheric sciences and meteorology communities. But improvements in the generative capabilities would likely come at the expense of the representation learning and the VAEs diagnosis of the physics of convection. We believe these preliminary physical intuitions achieved via representation learning represent a promising avenue for the broader application of generative modeling for advancing the field of atmospheric dynamics [1, 11] and warrant further investigation to understand their full potential.

ACKNOWLEDGMENTS

The authors thank Liran Peng for assistance with data generation as well as Karthik Mukkavilli and Ruihan Yang for helpful discussions that advanced this research project. GM is grateful to the MAPS Program and NSF grant 1633631 for funding support. TB is supported by NSF grants OAC-1835769, OAC-1835863, and AGS-1734164. SM acknowledges funding from DARPA (HR001119S0038). MP is also supported by NSF grant OAC-1835863. The code for this study can be found at <https://github.com/gmoosers96/CBRAIN-CAM/tree/master/MAPS>

REFERENCES

- [1] Alexander A. Alemi, Ben Poole, Ian S. Fischer, Joshua V. Dillon, Rif A. Saurous, and Kevin Murphy. 2018. Fixing a Broken ELBO. In *ICML*.
- [2] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Józefowicz, and Samy Bengio. 2016. Generating Sentences from a Continuous Space. In *CoNLL*.
- [3] Adam Clark, William Gallus, and Tsing-Chang Chen. 2007. Comparison of the Diurnal Precipitation Cycle in Convection-Resolving and Non-Convection-Resolving Mesoscale Models. *Monthly Weather Review - MON WEATHER REV* 135 (10 2007). <https://doi.org/10.1175/MWR3467.1>
- [4] J. Cooley and John W. Tukey. 1965. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.* 19 (1965), 297–301.
- [5] Daan Crommelin and Wouter Edeling. 2020. Resampling with neural networks for stochastic parameterization in multiscale systems. arXiv:2004.01457 [math.NA]
- [6] Stephan Eismann, Stefan Bartsch, and Stefano Ermon. 2017. Shape optimization in laminar flow with a label-guided variational autoencoder. arXiv:1712.03599 [cs.CE]
- [7] David John Gagne II, Hannah M. Christensen, Aneesh C. Subramanian, and Adam H. Monahan. 2020. Machine Learning for Stochastic Parameterization: Generative Adversarial Networks in the Lorenz '96 Model. *Journal of Advances in Modeling Earth Systems* 12, 3 (2020), e2019MS001896. <https://doi.org/10.1029/2019MS001896> arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS001896 e2019MS001896 10.1029/2019MS001896.
- [8] P. Gentile, M. Pritchard, S. Rasp, G. Reinaudi, and G. Yacalis. 2018. Could Machine Learning Break the Convection Parameterization Deadlock? *Geophysical Research Letters* 45, 11 (2018), 5742–5751. <https://doi.org/10.1029/2018GL078202>
- [9] Virendra P. Ghatge, Bruce A. Albrecht, and Pavlos Kollias. 2010. Vertical velocity structure of nonprecipitating continental boundary layer stratocumulus clouds. *Journal of Geophysical Research: Atmospheres* 115, D13 (2010). <https://doi.org/10.1029/2009JD013091> arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2009JD013091
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [11] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*.
- [12] Huaibo Huang, Zhihang Li, Ran He, Zhenan Sun, and Tieniu Tan. 2018. IntroVAE: Intropective Variational Autoencoders for Photographic Image Synthesis. arXiv:1807.06358 [cs.LG]
- [13] Eric J. Jensen, Glenn Diskin, R. Paul Lawson, Sara Lance, T. Paul Bui, Dennis Hlavka, Matthew McGill, Leonhard Pfister, Owen B. Toon, and Rushan Gao. 2013. Ice nucleation and dehydration in the Tropical Tropopause Layer. *Proceedings of the National Academy of Sciences* 110, 6 (2013), 2041–2046. <https://doi.org/10.1073/pnas.1217104110> arXiv:https://www.pnas.org/content/110/6/2041.full.pdf
- [14] Todd R. Jones, David A. Randall, and Mark D. Branson. 2019. Multiple-Instance Superparameterization: 1. Concept, and Predictability of Precipitation. *Journal of Advances in Modeling Earth Systems* 11, 11 (2019), 3497–3520. <https://doi.org/10.1029/2019MS001610> arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS001610
- [15] Heike Kalesse and Pavlos Kollias. 2013. Climatology of High Cloud Dynamics Using Profiling ARM Doppler Radar Observations. *Journal of Climate* 26, 17 (2013), 6340–6359. <https://doi.org/10.1175/JCLI-D-12-00695.1> arXiv:https://doi.org/10.1175/JCLI-D-12-00695.1
- [16] M. Khairoutdinov and Y. Kogan. 1999. A Large Eddy Simulation Model with Explicit Microphysics: Validation against Aircraft Observations of a Stratocumulus-Topped Boundary Layer. *Journal of the Atmospheric Sciences* 56 (1999), 2115–2131.
- [17] Marat Khairoutdinov and David Randall. 2003. Cloud Resolving Modeling of the ARM Summer 1997 IOP: Model Formulation, Results, Uncertainties, and Sensitivities. *Journal of The Atmospheric Sciences - J ATMOS SCI* 60 (02 2003), 607–625. [https://doi.org/10.1175/1520-0469\(2003\)060<0607:CRMOTA>2.0.CO;2](https://doi.org/10.1175/1520-0469(2003)060<0607:CRMOTA>2.0.CO;2)
- [18] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. *CoRR* abs/1312.6114 (2014).
- [19] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. 2017. CIFAR-10 (Canadian Institute for Advanced Research). *University of Toronto* (2017). <http://www.cs.toronto.edu/~kriz/cifar.html>
- [20] Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. 2018. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications* 9, 1 (Nov 2018). <https://doi.org/10.1038/s41467-018-07210-0>
- [21] Brian Medeiros, Bjorn Stevens, and Sandrine Bony. 2015. Using aquaplanets to understand the robust responses of comprehensive climate models to forcing. *Climate Dynamics* 44, 7 (2015), 1957–1977. <https://doi.org/10.1007/s00382-014-2138-0>
- [22] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. 2017. Adversarial Variational Bayes: Unifying Variational Autoencoders and Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) (ICML '17). JMLR.org, 2391–2400.
- [23] Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. The MIT Press.
- [24] Sebastian K. Müller, Elisa Manzini, Marco Giorgetta, Kaoru Sato, and Tomoe Nasuno. 2018. Convectively Generated Gravity Waves in High Resolution Models of Tropical Dynamics. *Journal of Advances in Modeling Earth Systems* 10, 10 (2018), 2564–2588. <https://doi.org/10.1029/2018MS001390> arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018MS001390
- [25] Paul A. O’Gorman and John G. Dwyer. 2018. Using Machine Learning to Parameterize Moist Convection: Potential for Modeling of Climate, Climate Change, and Extreme Events. *Journal of Advances in Modeling Earth Systems* 10, 10 (2018), 2548–2563. <https://doi.org/10.1029/2018MS001351>
- [26] Karl Pearson. 1901. LIII. On lines and planes of closest fit to systems of points in space. <https://doi.org/10.1080/14786440109462720>
- [27] David Randall, Marat Khairoutdinov, Akio Arakawa, and Wojciech Grabowski. 2003. Breaking the Cloud Parameterization Deadlock. *Bulletin of the American Meteorological Society* 84, 11 (2003), 1547–1564. <https://doi.org/10.1175/BAMS-84-11-1547> arXiv:https://doi.org/10.1175/BAMS-84-11-1547
- [28] Stephan Rasp, Michael S. Pritchard, and Pierre Gentile. 2018. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences* 115, 39 (2018), 9684–9689. <https://doi.org/10.1073/pnas.1810286115> arXiv:https://www.pnas.org/content/115/39/9684.full.pdf
- [29] Tapio Schneider, João Teixeira, Christopher Bretherton, Florent Briant, Kyle Pressel, Christoph Schär, and A.P. Siebesma. 2017. Climate goals and computing the future of clouds. *Nature Climate Change* 7 (01 2017), 3–5. <https://doi.org/10.1038/nclimate3190>
- [30] Panos Stinis, Tobias Hagge, Alexandre Tartakovsky, and Enoch Yeung. 2018. Enforcing constraints for interpolation and extrapolation in Generative Adversarial Networks. *J. Comput. Phys.* (03 2018). <https://doi.org/10.1016/j.jcp.2019.07.042>
- [31] Xavier-Andoni Tibau, Christian Requena-Mesa, Christian Reimers, Joachim Denzler, Veronika Eyring, Markus Reichstein, and Jakob Runge. 2018. SupernoVAE: VAE based Kernel-PCA for Analysis of Spatio-Temporal Earth Data. In *International Workshop on Climate Informatics (CI)*. 73–76. <https://doi.org/10.5065/D6BZ64XQ>
- [32] Natalia Tilinina, Mikhail Krinitskiy, Yulia Zyulyaeva, and Sergey Gulev. 2019. Clustering of the Polar Vortex states using deep convolutional neural networks. In *EGU General Assembly Conference Abstracts (EGU General Assembly Conference Abstracts)*. Article 17539, 17539 pages.

- [33] Jin-Long Wu, Karthik Kashinath, Adrian Albert, Dragos Chirila, Prabhat, and Heng Xiao. 2020. Enforcing statistical constraints in generative adversarial networks for modeling chaotic dynamical systems. *J. Comput. Phys.* 406 (2020), 109209. <https://doi.org/10.1016/j.jcp.2019.109209>
- [34] Yuhuai Wu, Yuri Burda, Ruslan Salakhutdinov, and Roger Grosse. 2017. On the Quantitative Analysis of Decoder-Based Generative Models. arXiv:1611.04273 [cs.LG]
- [35] Liu Yang, Dongkun Zhang, and George Em Karniadakis. 2020. Physics-Informed Generative Adversarial Networks for Stochastic Differential Equations. *SIAM J. Scientific Computing* 42 (2020), A292–A317.
- [36] Zeng Yang, Jin-Long Wu, and Heng Xiao. 2019. Enforcing Deterministic Constraints on Generative Adversarial Networks for Emulating Physical Systems. arXiv:1911.06671 [physics.comp-ph]