Modeling Cloud Reflectance Fields Using Conditional Generative Adversarial Networks

Victor Schmidt, Mustafa Alghali, Kris Sankaran, Tianle Yuan, Yoshua Bengio.

ICLR-CCAI 2020

All code and hyperparameters may be found at https://github.com/krisrs1128/clouds_dist

(1) Motivation

Global Climate Models (GCMs)

- GCMs had huge success in simulating the earth's weather, energy balance, and predicting possible changes in climate[1] including but not limited to:

changes in precipitation*

increases in temperatures**

acceleration in glacial melting***

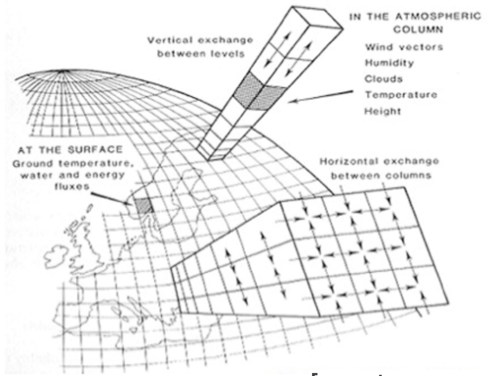

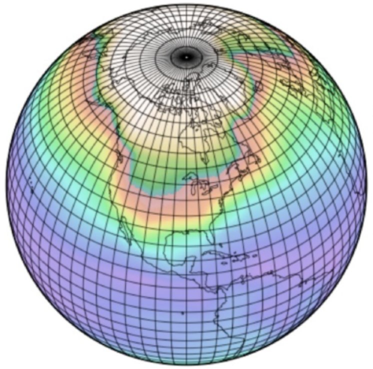

[Henderson and Sellers, 1985]

- One of the key physical principles these models rely on is the Earth's energy balance[2]

* USGS water science school ** Future impacts of climate change on forests *** scientificamerican.com

Clouds modeling and earth's energy balance

- Clouds play an important role in earth's energy balance as they both reflect energy coming to the Earth and the infrared radiations it emits.[3]

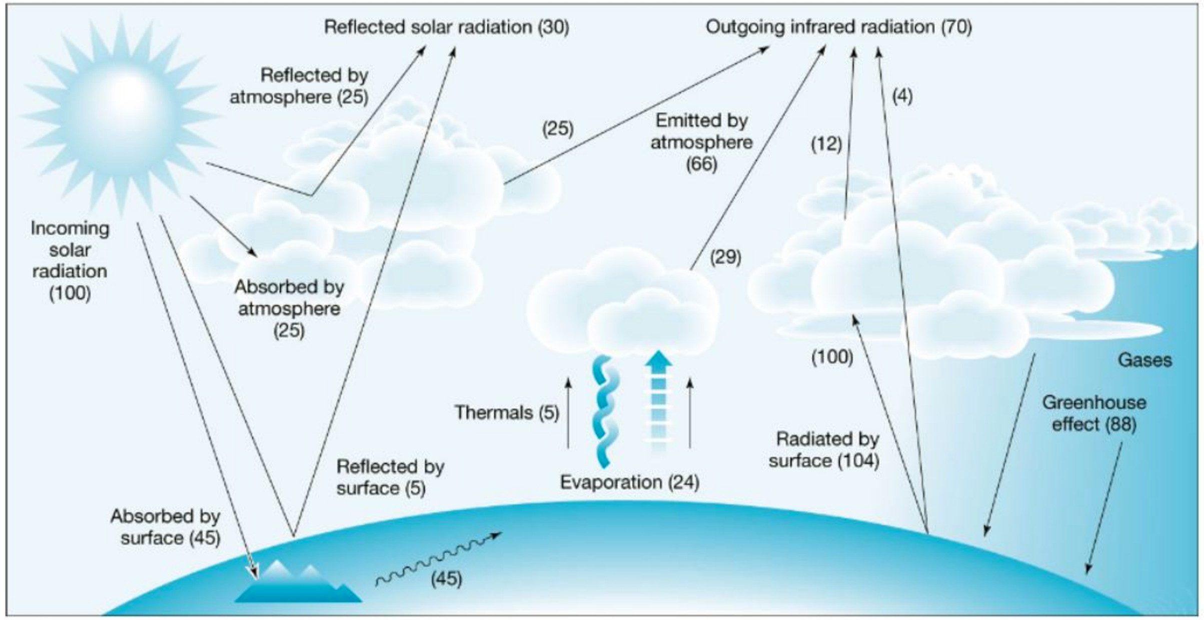

[Schneider, Stephen H. "Climate modeling." *Scientific American* 256.5 (1987): 72-T9]

- However, as physical processes at play in cloud composition and evolution typically range from 10^{-6} to 10^{6} m, direct simulation of their behavior can consume up to 20% of a GCM's computations.[4, 5, 6]

- Modeling clouds accurately using GCMs is challenging and expensive.

Cloud modeling computational complexity

Various efforts have tried to address this challenge such as:

- Incorporate more domain knowledge

- super-parameterization (modeling sub-grids)

✓ Machine learning (model sub-grid using meteorological variables) [7, 8, 9, 10]

(2) Approach

Narrowing down the clouds modeling challenge

🎯 In our approach we propose modeling Cloud Reflectance Fields (CRFs) using conditional Generative Adversarial Networks (GANs)

- We suggest using the generated CRFs as a proxy from which we can extract important cloud parameters such as optical depth and integrate these parameters into GCMs (it is not an alternative to GCMs)

- We believe our approach is a step towards building a data-driven framework that can reduce the computational complexity in traditional cloud modeling techniques.

Approach: overview

- We use GAN to generate cloud reflectance fields conditioned on meteorological variables, taking the climate chaotic nature into consideration.

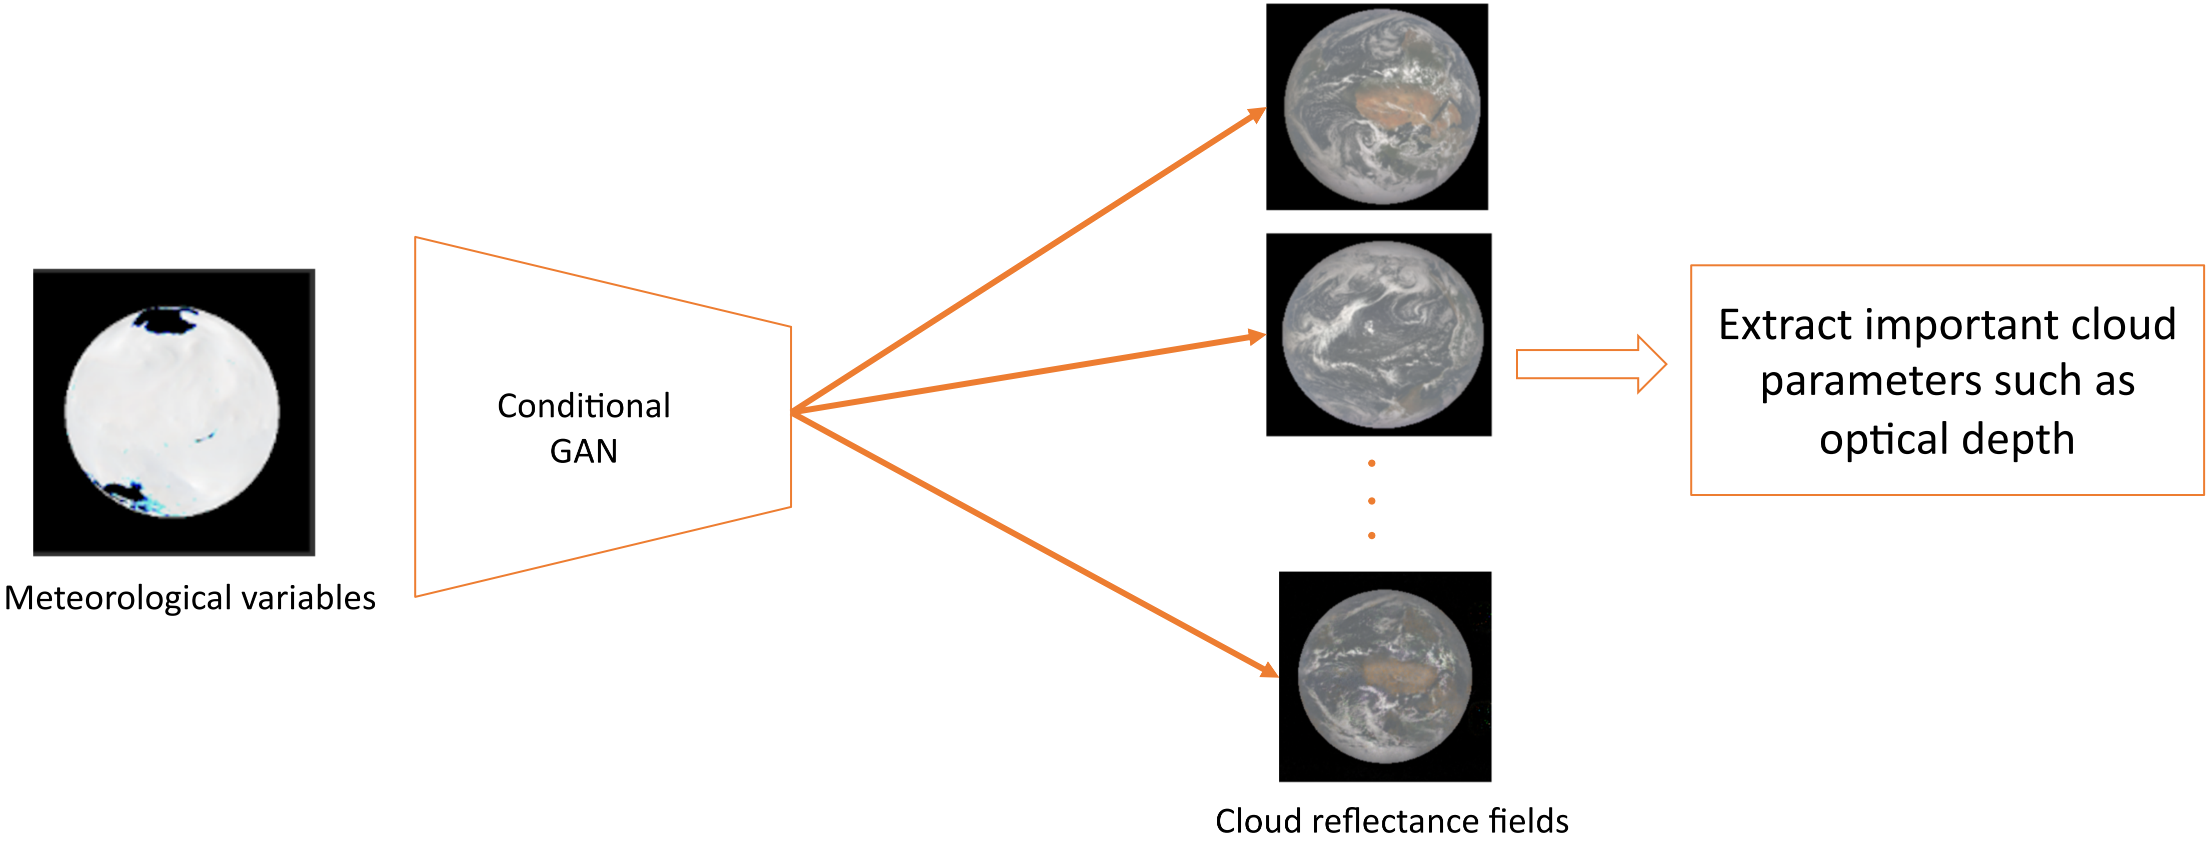

Approach: Data

- Training data: 3100 aligned sample pairs $X = \{m_i, r_i\}$
- Independent variable (m_i) 🌡️ : is a 44 × 256 × 256 matrix, representing 42 measurements from NASA's MERRA-2[19] along with longitude and latitude to account for the Earth's movement relative to the satellite.
- Dependent variable (r_i) 🌏: is a 3 × 256 × 256 matrix representing each location's reflectance at RGB wavelengths (680, 550 and 450 nm) as measured by the Aqua dataset [20].

(3) Methodology

Architecture: Generator

- **U-Net generator** [11]
  - Skip connections help localization
  - reduce the need for larger training set

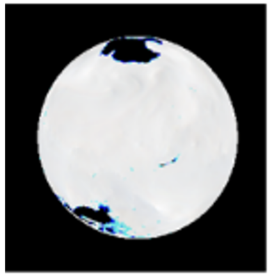

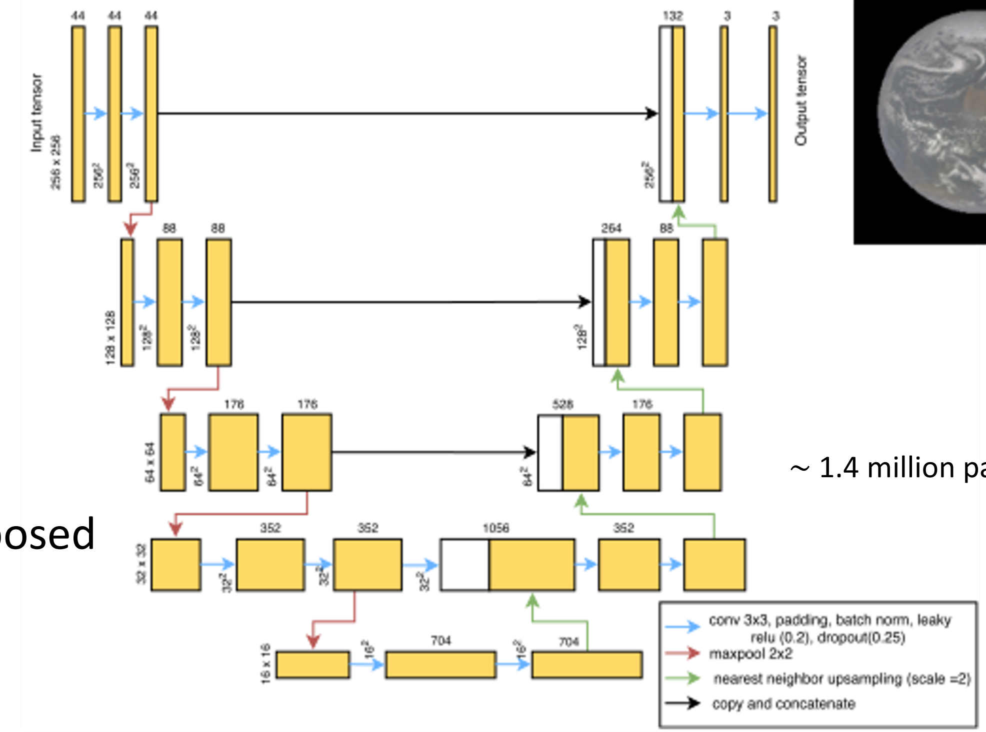

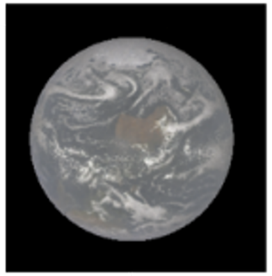

~ 1.4 million parameters

- **Checkerboard artifacts** [12]
  - Upsampling followed by a convolution instead of transposed convolution

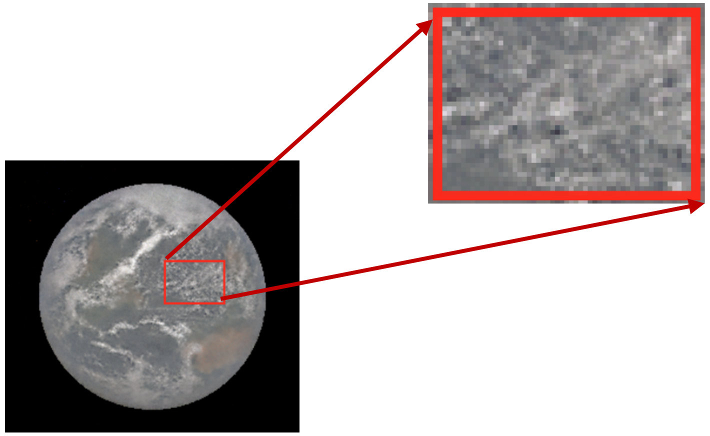

Architecture: Discriminator

- **Multi-scale discriminator** [13]
  - Better guide for the generator both in the scale of global context and finer details in the image.

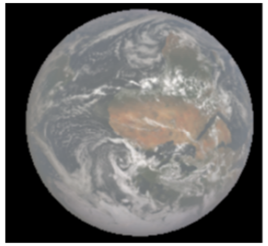

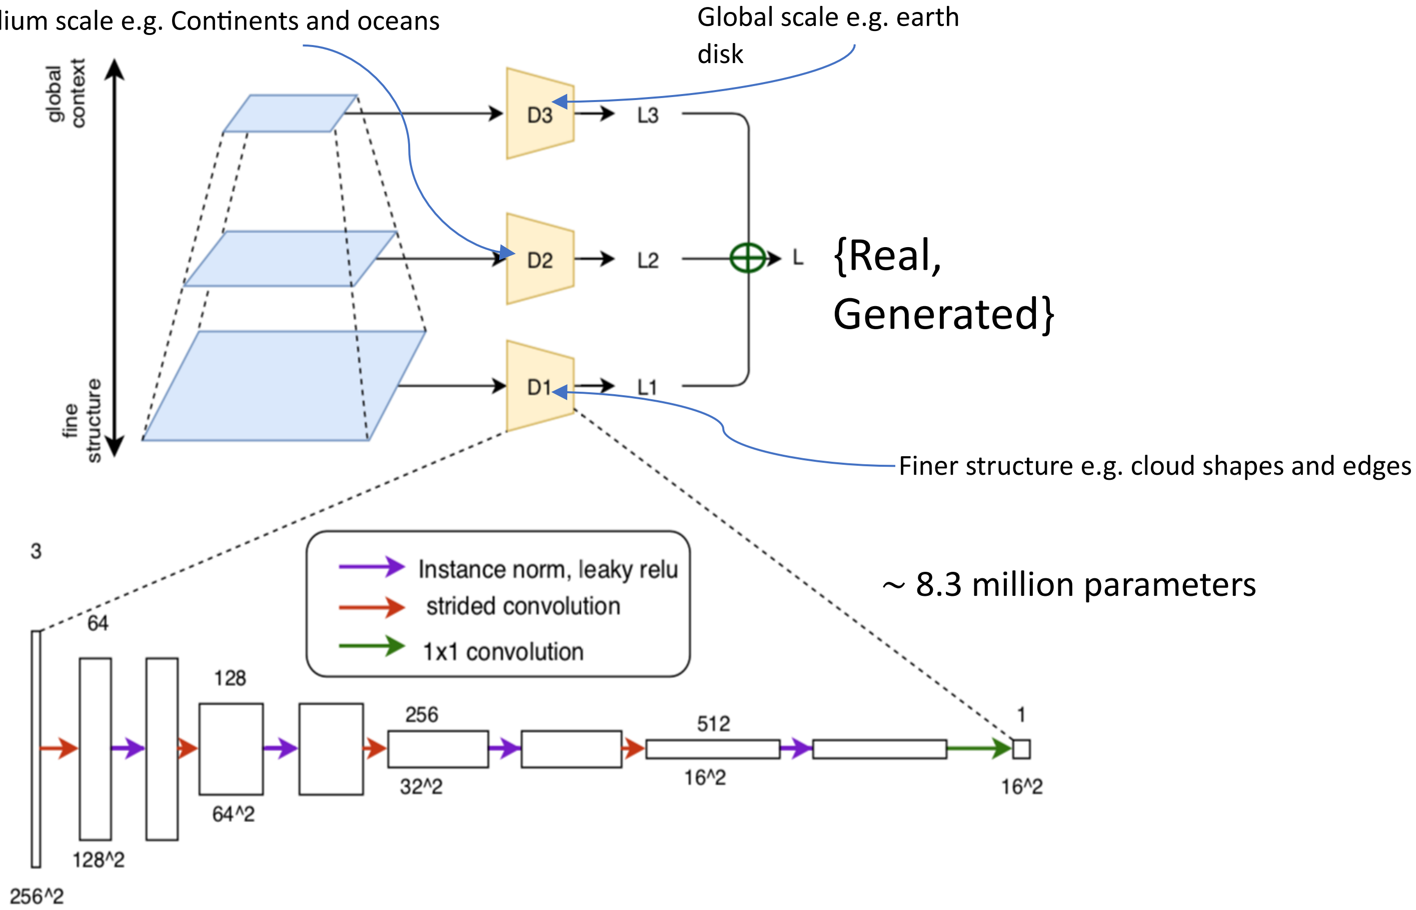

~ 8.3 million parameters

Training objective

Total GAN loss

$\lambda_1 *$ Non saturating adversarial loss

$\lambda_2 *$ Matching loss [16]

Less blurry output than L_2 loss

Least square loss (LSGAN) [14]

$$\min_G \mathbb{E}_m \left[\left(D_k(G(m)) - c \right)^2 \right], \forall k = 1, 2, 3$$

Hinge loss [15]

$$\min_G \mathbb{E}_m \left[-D_k(G(m)) \right], \forall k = 1, 2, 3$$

Challenges: Optimization

- Adam/SGD
- Extra_SGD [17]

✔ **Extra-Adam** [17]

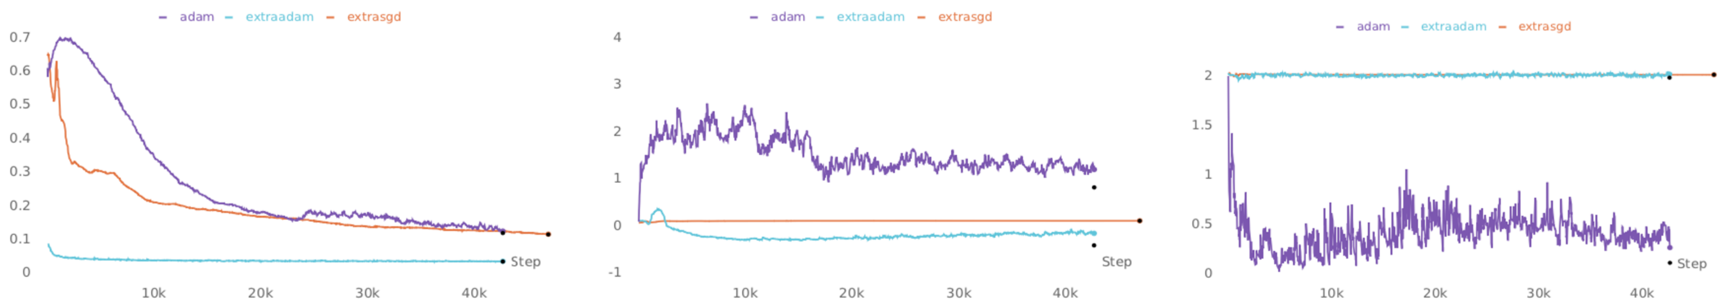

(a) L1 matching loss (b) Generator adversarial loss (c) Discriminator loss

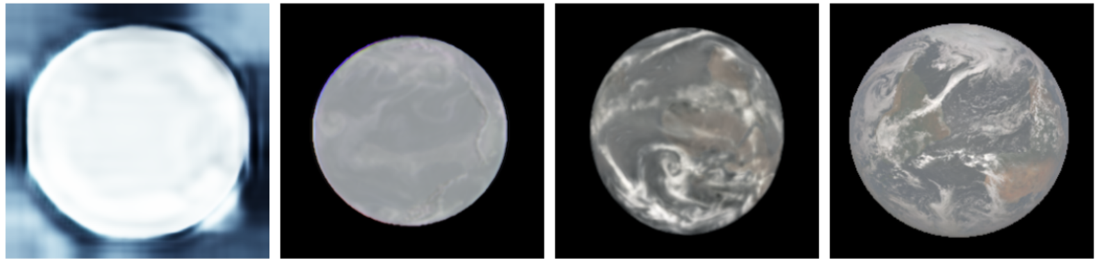

(d) Adam (e) ExtraSGD (f) ExtraAdam (g) Real earth

see code at https://github.com/GauthierGidel/Variational-Inequality-GAN

Challenges: Regression vs. hallucinated features

$\frac{\lambda_1}{\lambda_2} \in$

$[0.5, 0]$

- behave more like supervised regression problem

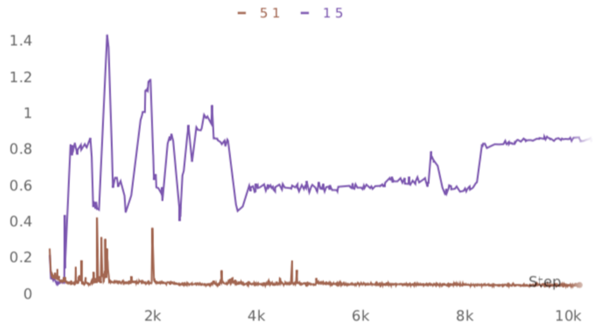

(a) L^1 matching loss

$[1, 10]$

- more freedom to explore the distribution of interest
- hallucinate features on cost of low frequency features

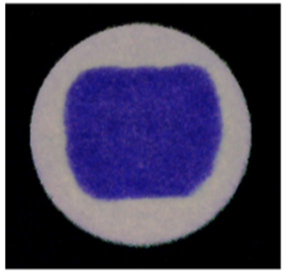

Challenges: Sharpness of generated images

- Prematurely saturated learning (Nash equilibrium) [18]

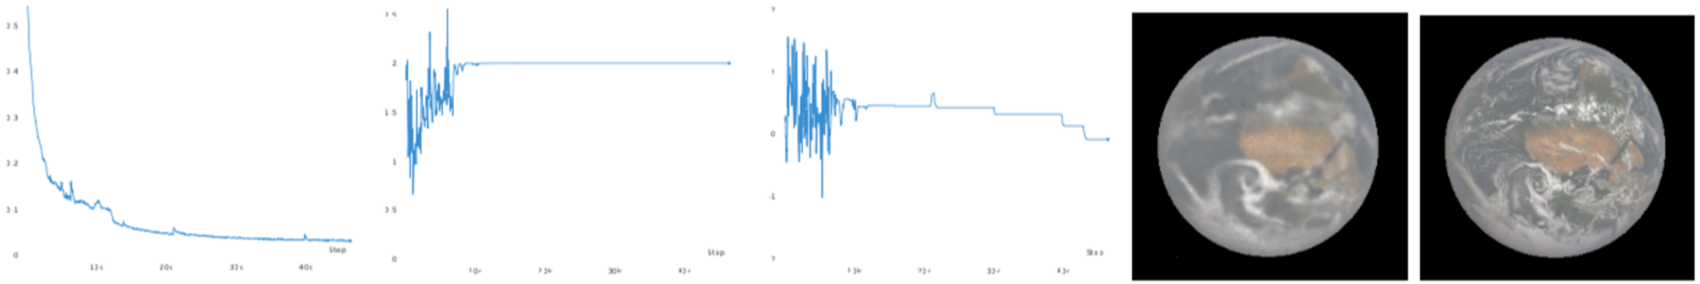

(a) L^1 matching loss (b) Generator loss (c) Discriminator loss (d) fake (e) real

- Carefully choose the discriminator learning rate! 🎯

(4) Results

Visual Analysis

- Generated images look difficult to distinguish from true samples with average L_2 distance ~ 0.027 on validation set.

- Validation set is set to 5 samples that are selected manually to capture different regions of the rotating earth.

- Generate 15 samples in total: 3 for each validation sample.

Generated (left) Real (right) Generated (left) Real (right)

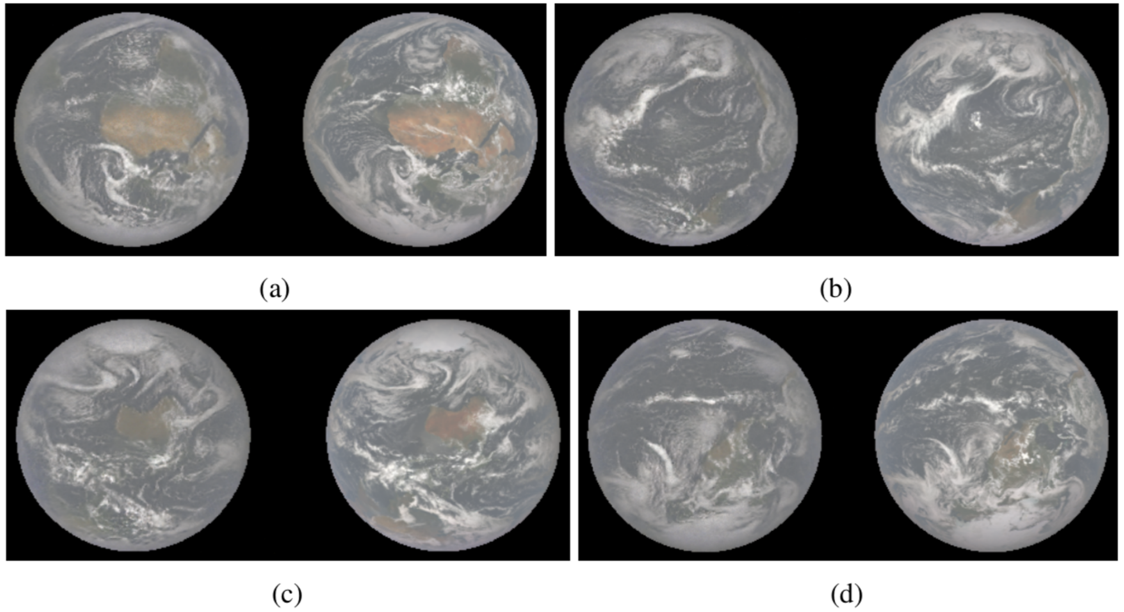

Model inference on never seen examples

Visual Analysis: Quantifying ensemble diversity

- For each ensemble generation we calculate:
  - Pixel-wise mean
  - Standard deviation
  - Inter-quartile range (IQR)

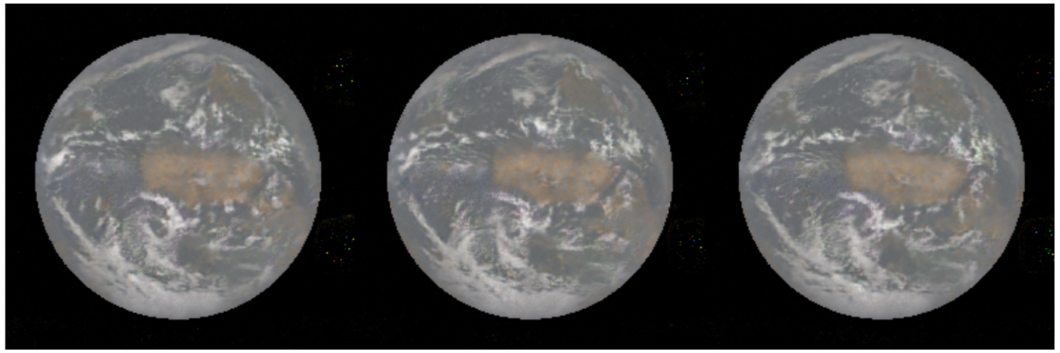

Ensemble generation conditioned on the same input

- Tradeoff (generation quality ↔ generation diversity)

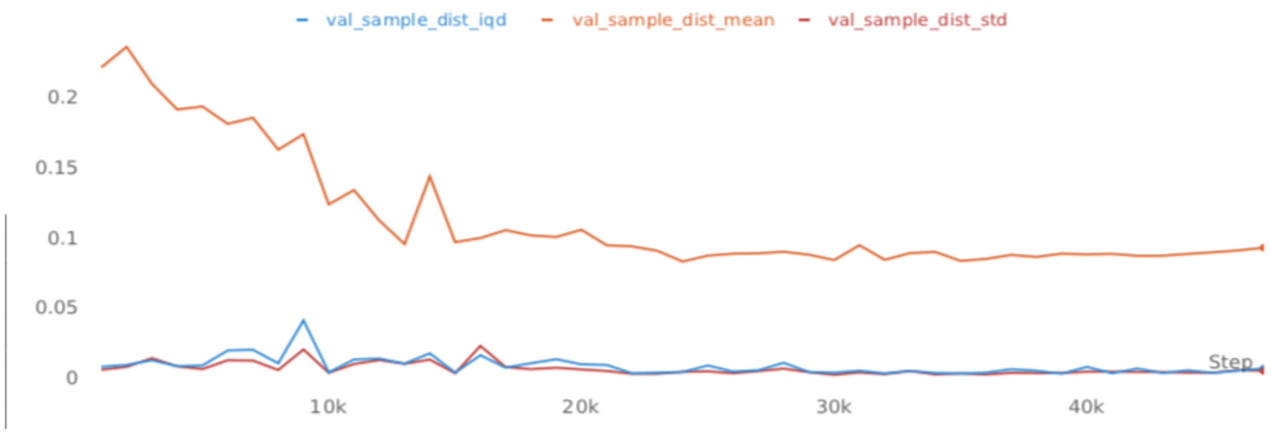

Spectral Analysis

- Visual inspection is an expensive, cumbersome, and subjective measure!
- Spectral analysis:
  - ✔Similar DFT distributions but there is still room for improvement
  - ✔Very small average L2 loss of 0.006 per frequency component.

$$F(k,l) = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} f(i,j)^{-i2\pi(\frac{ki}{N}+\frac{lj}{N})}$$

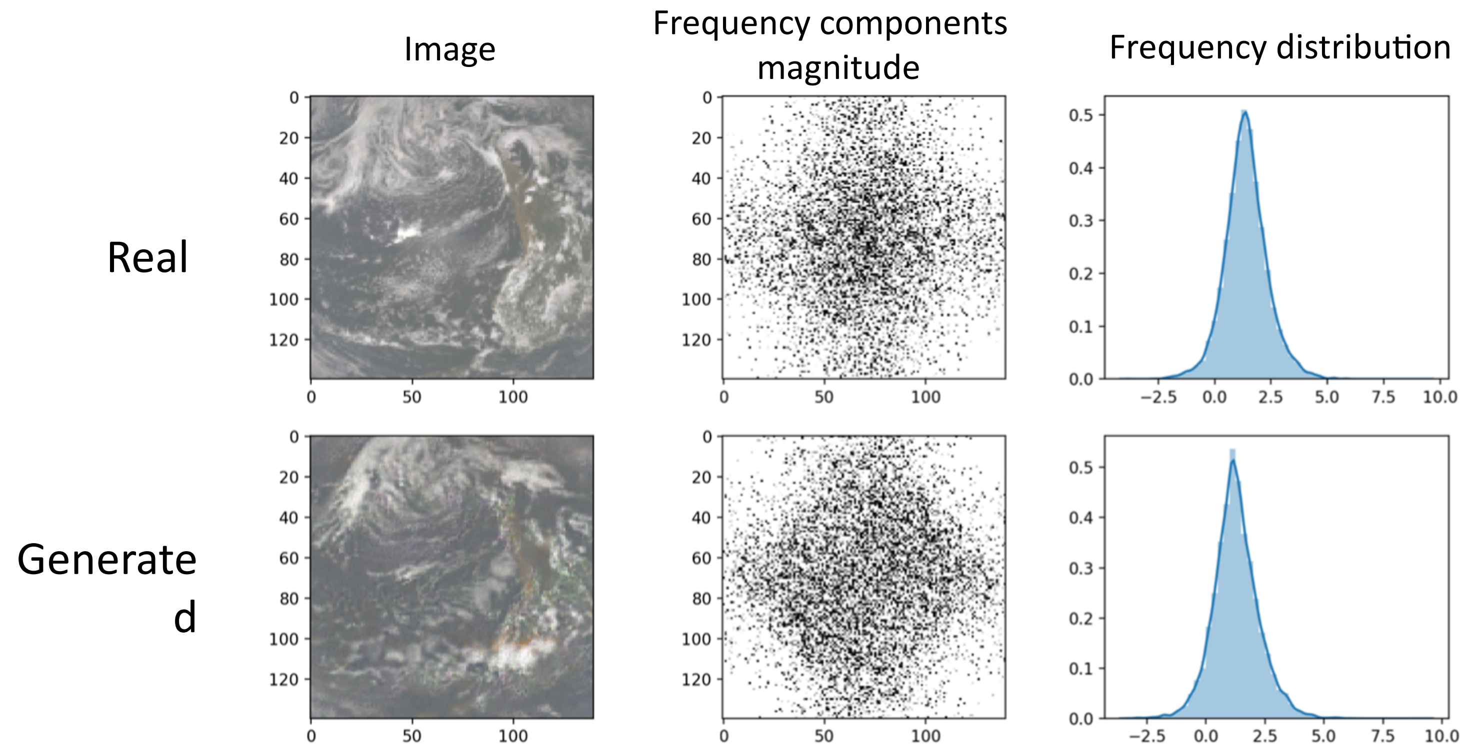

What’s next?

- Blurriness and small size checkerboard artifacts:

  ❑ More training samples

  ❑ More hyperparameter tuning → avoid prematurely saturated learning.

  ❑ Longer training

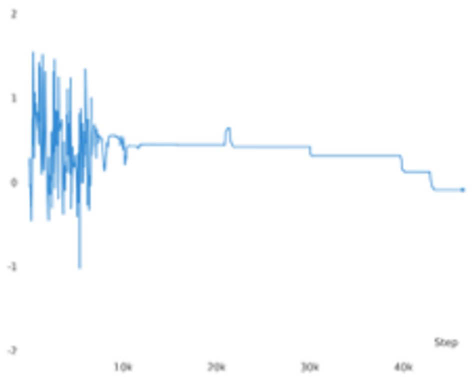

What's next?

- Exploit temporal structure:
  - Add date and time as extra labels to the input variable.
  - Using nested temporal cross validation to predict possible changes in cloud distribution over time.

- Increase the diversity in the generated ensembles.
  - Incorporate input noise channels as an extra source of stochasticity
  - Address mode collapse by using decaying λ2

λ_2 =exp(-t)

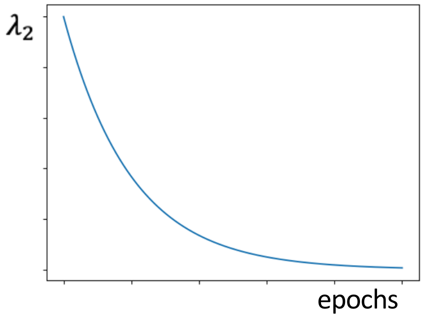

- Modeling low clouds a key source of uncertainty in our ability to project future climate changes [21]

Appendix A: Data

Table 1: Description of input components

Name	Description	Number of channels
U, V	Wind components in 10 atmospheric levels	20
T	Temperature in 10 atmospheric levels	10
RH	Relative-humidity in 10 atmospheric levels	10
SA	Scattering angle	1
TS	Surface Temperature	1
Lat, Long	Latitude and Longitude	2

Appendix B: Data processing

- Sensor noise Winsorization → clip CRFs to the 95^{th} percentile.

- Standardization

- Avoid introducing unnecessary bias in the data distribution by the values outside the earth disk
  - Reduce them by zooming (crop & then resize using 2D nearest neighbor)
  - Replace other remaining values with -3 (mean - 3x standard deviation)

- Use running statistics → mitigate shortage of GPU memory budget

- Use 12 data loader workers → speed up the data loading process 6x

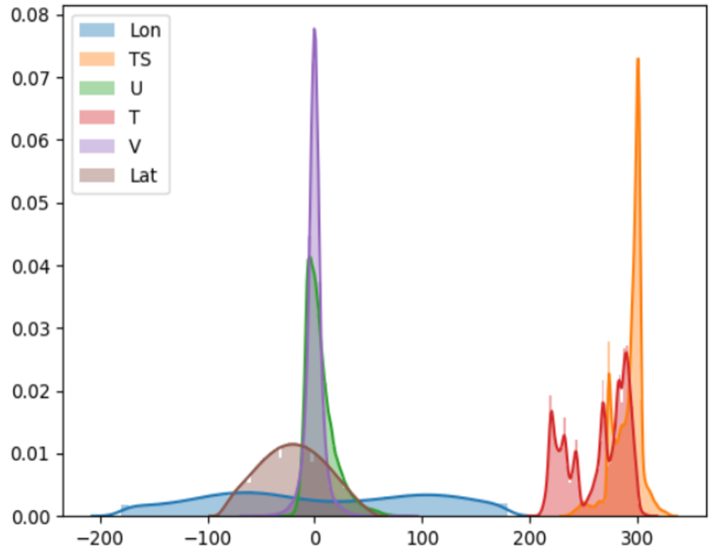

Appendix C: Hyperparameters

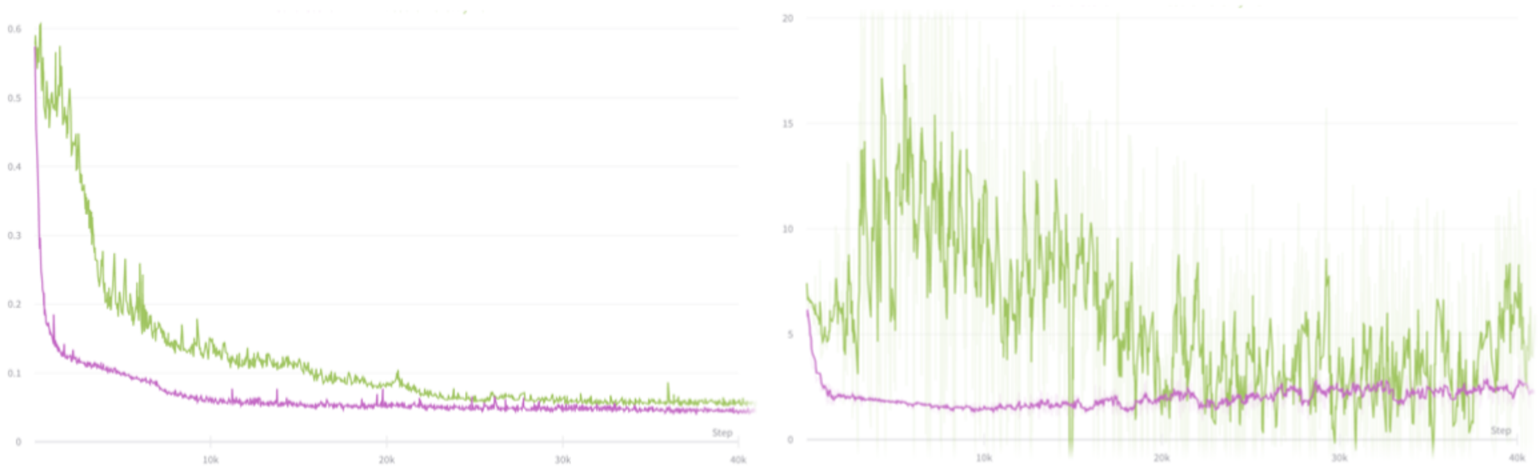

(a) L1 matching loss (b) Total weighted generator loss

Figure 4: Comparison between the hing loss shown in green and the least squares loss shown in purple on model training stability and convergence, we can observe that the latter is performing better both in optimization of the L1 loss and the total weighted generator loss

References:

[1] Thomas F Stocker, Dahe Qin, Gian-Kasper Plattner, Melinda Tignor, Simon K Allen, Judith Boschung, Alexander Nauels, Yu Xia, Vincent Bex, Pauline M Midgley, et al. Climate change 2013: The physical science basis. *Contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change*, 1535, 2013.

[2] Gerald R North, Robert F Cahalan, and James A Coakley Jr. Energy balance climate models. *Reviews of Geophysics*, 19(1):91–121, 1981.

[3] VLRD Ramanathan, RD Cess, EF Harrison, P Minnis, BR Barkstrom, E Ahmad, and D Hart- mann. Cloud-radiative forcing and climate: Results from the earth radiation budget experiment. *Science*, 243(4887):57–63, 1989.

[4] Akio Arakawa. The cumulus parameterization problem: Past, present, and future. *Journal of Climate*, 17(13):2493–2525, 2004.

[5] Christopher S Bretherton. Insights into low-latitude cloud feedbacks from high-resolution models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2054):20140415, 2015.

[6] Tapio Schneider, João Teixeira, Christopher S Bretherton, Florent Brient, Kyle G Pressel, Christoph Schär, and A Pier Siebesma. Climate goals and computing the future of clouds. *Nature Climate Change*, 7(1):3–5, 2017.

References:

[7] Noah D Brenowitz and Christopher S Bretherton. Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12):6289–6298, 2018.

[8] Stephan Rasp, Michael S Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39): 9684–9689, 2018.

[9] Paul A O'Gorman and John G Dwyer: 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):2548–2563, 2018.

[10] T. Yuan, H. Song, D. Hall, V. Schmidt, K. Sankaran, and Y. Bengio. Artificial intelligence based cloud distributor (ai-cd): probing clouds with generative adversarial networks. *AGU Fall Meeting 2019*, 2019.

[11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

[12] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 1(10):e3, 2016.

References:

[13] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Pro- ceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807, 2018.

[14] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2794–2802, 2017.

[15] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.

[16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.

[17] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018.

[18] Salimans, Tim, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. "Improved techniques for training gans." In *Advances in neural information processing systems*, pp. 2234-2242. 2016

References:

[19] Ronald Gelaro, Will McCarty, Max J Suárez, Ricardo Todling, Andrea Molod, Lawrence Takacs, Cynthia A Randles, Anton Darmenov, Michael G Bosilovich, Rolf Reichle, et al. The modern-era retrospective analysis for research and applications, version 2 (merra-2). *Journal of Climate*, 30(14):5419–5454, 2017.

[20] S Platnick, KG Meyer, MD King, G Wind, N Amarasinghe, B Marchant, GT Arnold, Z Zhang, PA Hubanks, RE Holz, et al. The modis cloud optical and microphysical products: Collection 6 updates and examples from terra and aqua, ieee t. geosci. remote, 55, 502–525, 2017.

[21] Bony, Sandrine, and Jean-Louis Dufresne. "Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models." *Geophysical Research Letters* 32, no. 20 (2005).