

# Cumulo: A Dataset for Learning Cloud Classes

Valentina Zantedeschi<sup>\*1</sup>, Fabrizio Falasca<sup>†2</sup>, Alyson Douglas<sup>3</sup>, Richard Strange<sup>4</sup>,  
Matt J. Kusner<sup>‡5,6</sup>, and Duncan Watson-Parris<sup>§4</sup>

<sup>1</sup>GE Global Research, San Ramon, USA

<sup>2</sup>Georgia Institute of Technology, Atlanta, USA

<sup>3</sup>University of Wisconsin, Madison, USA

<sup>4</sup>University of Oxford, Oxford, UK

<sup>5</sup>The Alan Turing Institute, London, UK

<sup>6</sup>University College London, London, UK

## Abstract

One of the greatest sources of uncertainty in future climate projections comes from limitations in modelling clouds and in understanding how different cloud types interact with the climate system. A key first step in reducing this uncertainty is to accurately classify cloud types at high spatial and temporal resolution. In this paper, we introduce CUMULO, a benchmark dataset for training and evaluating global cloud classification models. It consists of one year of 1km resolution MODIS hyperspectral imagery merged with pixel-width ‘tracks’ of CloudSat cloud labels. Bringing these complementary datasets together is a crucial first step, enabling the Machine-Learning community to develop innovative new techniques which could greatly benefit the Climate community. To showcase CUMULO, we provide baseline performance analysis using an invertible flow generative model (IResNet), which further allows us to discover new sub-classes for a given cloud class by exploring the latent space. To compare methods, we introduce a set of evaluation criteria, to identify models that are not only accurate, but also physically-realistic.

## 1 Cloud Classification is key for modelling Climate Change

Clouds play a crucial role in the climate system. They are the source of all precipitation and have a significant impact on the Earth’s radiative budget. Crucially, as any changes in clouds impact the environment; these changes feedback on cloud formation and behaviour. These feedbacks are a primary source of uncertainty for climate model projections [Knutti et al., 2017, Rotstayn and Collier, 2015, Stephens, 2005], as there is a limited understanding of the mechanisms and relationships between clouds, climate and global circulation [Bony et al., 2015]. It is for example unclear how warmer sea surface temperature will affect clouds and convective organization [Bony et al., 2015], or trigger possible climate transitions by cloud breakup [Schneider et al., 2019]. Comprehensive use of the vast observational data available is crucial to improve our understanding of these processes, and their representations in climate models.

Clouds can form and develop through several different pathways, depending on their environment and the convective energy available. It is common to categorise clouds into different types based on their properties to better analyse them. The International Satellite Cloud Climatology Project (ISCCP) dataset [Rossow and Schiffer, 1991] provides a global classification of clouds at a 10km resolution, based on a network of

---

\*vzantedeschi@gmail.com

†fabrifalasca@gatech.edu

‡mkusner@turing.ac.uk

§duncan.watson-parris@physics.ox.ac.uk

geostationary meteorological satellites. Satellite based observations of clouds can be made using either passive imagery or active radar instruments. While high-resolution hyperspectral imagery is available from both polar orbiting, and geostationary satellites, providing excellent coverage at high temporal resolution, specific cloud properties (such as their exact height and droplet size distribution) must be inferred indirectly. The ISCCP classification relies on a simple assessment of the relationship between the clouds’ inferred height and optical thickness [Rossow and Schiffer, 1991]. Conversely, the CloudSat cloud radar does provide direct measurements of clouds and their properties. This comes with a drawback, as it operates with a narrow swath and and repeat a cycle every 16 days (no global coverage at 1km resolution is present even after 16 days).

To overcome these limitations, in this paper we introduce CUMULO, a new dataset which combines the global 1km-resolution imagery of the Moderate Resolution Imaging Spectroradiometer (MODIS) with the accurately measured properties of the CloudSat products. It contains one year of 1354 x 2030 pixel hyperspectral images from MODIS combined with pixel-width ‘tracks’ of cloud labels from Cloudsat, corresponding to the eight World Meteorological Organization (WMO) genera (Fig. 1). While both datasets are publicly available, the extraction, cleaning and alignment of the data required specialist domain knowledge and extensive compute resources.

We apply a deep generative model architecture on one month of CUMULO, and present, for the first time to our knowledge, global high resolution spatiotemporal cloud classification derived from a combination of active and passive satellite sensors. We show that our results are physically reasonable in terms of locations of occurrences of the given classes and liquid water path distributions.

**Related Work.** Muhlbauer et al. [2014] classify MODIS into three types of mesoscale cellular convection using a 3-layer neural network. While these classifications are global, they only describe a particular climatology. Zhang et al. [2019] classify the geostationary Himawari-8 satellite into WMO cloud classes (from Cloudsat) using a random forest. Alternatively, our dataset provides global coverage (versus East Asia and Western Pacific for Himawari-8). Rasp et al. [2019] crowd-sourced human-level classifications of shallow trade clouds into four types: ‘sugar’, ‘flower’, ‘fish’, ‘gravel’. They evaluate an object detection method Lin et al. [2017] and a semantic segmentation method Ronneberger et al. [2015] to classify clouds into these types. Similar to Muhlbauer et al. [2014], this work only aims at a small aspect of cloud variability.

## 2 Cumulo: A global dataset for cloud classification

The proposed dataset contains 105,120 geolocated and hyperspectral images and provides a combination of channels from different sources (see Table 1): the selected radiance channels from MODIS AQUA Calibrated Radiances fully capture the physical properties needed for cloud classification and are meant to be used for training; MODIS AQUA Cloud Product channels are retrieved features describing cloud physical properties useful for validation; MODIS AQUA Cloud Mask detects the presence of a cloud; 2B-CLDCLASS-LIDAR provides the types of clouds spotted at different heights along the track of the satellites. Possible cloud types, corresponding to the eight WMO genera, are stratus (St), stratocumulus (Sc), cumulus (Cu, including cumulus congestus), nimbostratus (Ns), altocumulus (Ac), altostratus (As), deep convective (cumulonimbus, Dc), and high (Ci, cirrus and cirrostratus). Refer to Table 2 for a description of the classes.

Notice that some channels are available only at daytime, because they rely directly or indirectly on daylight radiances. In general, missing values due to artefacts were filled with the nearest (in time and space) available values.

More precisely, each satellite image (swath) is acquired at a given time  $t$  (one swath every five minutes) and at a given location  $l$  (each pixel is associated with a latitude-longitude pair):  $\{\mathbf{S}^{t,l}\}_{t=1,\dots,T;l=1,\dots,L}$ . In the following, we denote by

- $\mathbf{X}^{t,l} \in \mathbb{R}^{2030 \times 1354 \times 13}$  the thirteen training channels coming from MODIS AQUA Calibrated Radiances <sup>1</sup>;
- $\mathbf{V}^{t,l} \in \mathbb{R}^{2030 \times 1354 \times 8}$  the eight validation channels coming from MODIS AQUA Cloud Product <sup>2</sup> which

<sup>1</sup><https://modis.gsfc.nasa.gov/data/dataproduct/mod02.php>

<sup>2</sup><https://modis.gsfc.nasa.gov/data/dataproduct/mod06.php>

Table 1: Channel descriptions

Source	Name/Description	Index	Primary Use	Availability
MODIS	shortwave visible (red)	1	land/shadow/cloud/aerosols boundaries	daytime
	shortwave near infrared	2	land/shadow/cloud/aerosols boundaries	daytime
	longwave thermal-infrared	20-23	surface/cloud temperature	always
	shortwave near infrared	26	Cirrus clouds water vapor	daytime
	longwave thermal-infrared	27	water vapor	always
	longwave thermal-infrared	29	cloud properties	always
	longwave thermal-infrared	33-36	cloud top altitude	always
MODIS Cloud Mask	cloud mask		cloud detection	always
2B-CLDCLASS-LIDAR	cloud layer type		target classes	always (limited coverage)
MODIS Cloud Product	liquid water path		physical validation	always
	cloud optical thickness		physical validation	daytime
	cloud effective radius		physical validation	daytime
	cloud particle phase		physical validation	daytime
	cloud top pressure		physical validation	always
	cloud top height		physical validation	always
	cloud top temperature		physical validation	always
	cloud effective emissivity		physical validation	always
	surface temperature		physical validation	always

Table 2: Class descriptions

Index	WMO Name	Characteristics				Proportion
		thickness	base height	liquid water path	rain	
0	Cirrus and cirrostratus (Ci)	moderate	> 7.0 km	0.	none	30.68%
1	Altostratus (As)	moderate	2.0-7.0 km	~ 0.	none	16.02%
2	Altostratus (Ac)	shallow/moderate	2.0-7.0 km	> 0.	virga possible	9.53%
3	Stratus (St)	shallow	0-2.0 km	> 0.	none/slight	1.84%
4	Stratocumulus (Sc)	shallow	0-2.0 km	> 0.	drizzle/snow possible	27.53%
5	Cumulus (Cu)	shallow/moderate	0-3.0 km	> 0.	drizzle/snow possible	6.02%
6	Nimbostratus (Ns)	thick	0-4.0 km	> 0.	prolonged rain/snow	7.40%
7	Deep Convection (Dc)	thick	0-3.0 km	> 0.	intense rain/hail	0.96%

provide physical and radiative cloud properties obtained by combining infrared emission and solar reflectance techniques applied on MODIS original bands;


- $\mathbf{C}^{t,l} \in \mathbb{R}^{2030 \times 1354}$  the cloud mask derived from MODIS Cloud Mask <sup>3</sup>, marking as 1 the certainly cloudy pixels and as 0 any other pixel;
- $\mathbf{L}^{t,l} \in \mathbb{R}^{2030 \times 1354 \times 8}$  the label mask derived from 2B-CLDCLASS-LIDAR <sup>4</sup>, indicating for each pixel the number of occurrences of the 8 cloud classes (refer to Table 2) that can be identified at different heights (10 layers maximum).

Figure 1 shows a snapshot of CUMULO for one day for the all globe (Fig. 1a) and for one swath (Fig. 1b), along with its MODIS cloud mask (Fig. 1c).


Overall, CUMULO provides a comprehensive set of features both for identifying cloud types and for validating any finding, along with ready-to-use and accurate cloud annotations at high spatial resolution. From a Machine Learning perspective, CUMULO presents several challenges: supervision is available only for 1 every 1354 pixels (weakly-labelled data), pixels can be annotated with multiple types of clouds (multi-labelled data), and many cloud classes, such as deep convective clouds, are underrepresented (class imbalance).

<sup>3</sup><https://modis.gsfc.nasa.gov/data/dataproduct/mod35.php>

<sup>4</sup><http://www.cloudsat.cira.colostate.edu/data-products/level-2b/2b-cldclass-lidar>



(a) All swaths of one day, projected on the globe, with their label masks. Notice that swaths can overlap and that, for visualisation purposes, label tracks are magnified.



(b) The visible band (left) and the cloud mask (right) of a sample swath with its overlying label mask.

Figure 1: Visualization of CUMULO’s data coverage.

### 3 Applying Cloud Classification globally

In this section, we provide baseline performance analysis of one of the tasks that can be performed using CUMULO: semi-supervised classification of clouds at a global and daily scale.

#### 3.1 Methodology

Given the small number of labels, we find that performing this experiment at a pixel level allows us to achieve better classification performance than applying common semantic segmentation models (such as Papandreou et al., Li et al. [2018], Ronneberger et al. [2015]) that necessitate full label masks or annotations at the instance or image level. We consider tiles of 3x3 pixels extracted from the first month of data (January 2008) and predict a label for each tile. We use the most frequent cloud type identified within each tile as a target, when annotations are available. We have gathered two sets of tiles: a set of labelled tiles sampled around any annotated pixel on the satellite’s track; and a sample of unlabelled tiles randomly selected from the non-annotated regions. We have made use of the available cloud mask for restraining the classification to tiles that had a high probability of cloud cover. We have trained a hybrid Invertible Residual Network [Nalisnick et al., 2019] which allows us to (i) harness both labelled and unlabelled sets and (ii) learn a representation where the class distributions can be further subdivided into fine-grained classes. Indeed, cloud types are not limited to the well-studied WMO genera: being able to identify more species of clouds is an open-question

in the cloud community. Our model combines a deep generative flownet with a linear classifier, which is simultaneously trained by maximizing the joint log-likelihood over the tiles and their labels <sup>5</sup>. Classes are weighted in the objective to encourage a better classification of under-represented classes.

### 3.2 Formal problem setting

We segment each tensor  $\mathbf{X}^{t,l}$  into non-overlapping tiles of size 3x3 pixels, e.g.

$$\mathbf{X}^{t,l} = \begin{bmatrix} \mathbf{X}_{(0,0)}^{t,l} & \mathbf{X}_{(0,1)}^{t,l} & \cdots & \mathbf{X}_{(0,W)}^{t,l} \\ \mathbf{X}_{(1,0)}^{t,l} & \mathbf{X}_{(1,1)}^{t,l} & \cdots & \mathbf{X}_{(1,W)}^{t,l} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{X}_{(H,0)}^{t,l} & \mathbf{X}_{(H,1)}^{t,l} & \cdots & \mathbf{X}_{(H,W)}^{t,l} \end{bmatrix}, \quad (1)$$

and aim at learning a mapping  $f : \mathcal{X} = \mathbb{R}^{3,3,13} \rightarrow \mathcal{Y} = \{-1, \dots, 7\}$  from any tile  $\mathbf{X}_{(i,j)}^{t,l}$  to a class label  $\mathbf{y}_{(i,j)}^{t,l}$ . As target values, we retain the most frequent cloud type occurring within the label mask  $\mathbf{L}_{(i,j)}^{t,l}$  associated to a tile.

For learning, we recall that we make use of two sets of tiles of equal size  $K$ : a set of labeled tiles  $\mathcal{L} = \{(x_k, y_k) \in \mathcal{X} \times \mathcal{Y}\}_{k=1}^K$ , sampled around any annotated pixel on the satellite’s track, and a sample of unlabeled tiles  $\mathcal{U} = \{x_k \in \mathcal{X}\}_{k=1}^K$ , randomly selected from the non-annotated regions. We deploy a hybrid Invertible Residual Network [Nalisnick et al., 2019], parameterised by  $\theta$ , which allows us to learn a latent representation for modelling the true distribution of the data  $p_\theta(\mathbf{X})$  by the decomposition

$$p_\theta(\mathbf{X}) = \sum_{\mathbf{z}} p_\theta(\mathbf{X}|\mathbf{z})p(\mathbf{z}) \quad (2)$$

where  $\mathbf{z} \in \mathbb{R}^{2030,1354}$  is the latent representation of  $\mathbf{X}$  and  $p(\mathbf{z}) \sim \mathcal{N}(0, 1)$  its prior distribution.

The peculiarity of flownets, such as the hybrid IResNet, is that they consist of a series of mappings  $f = f_1 \circ f_2 \circ \dots \circ f_D$  that are invertible for any input  $x \in \mathcal{X}$ :  $\forall i = 1, \dots, D; f_i(x)^{-1} = x$ . Crucially, there is a direct and invertible relationship between each image  $\mathbf{X}$  and each latent point  $\mathbf{z}$ :

$$\mathbf{X} \xleftrightarrow{f_1} \mathbf{h}_1 \xleftrightarrow{f_2} \mathbf{h}_2 \dots \xleftrightarrow{f_D} \mathbf{z} \quad (3)$$

Thanks to these constraints on the architecture of the network, it is possible to optimise  $\theta$  by maximising the joint likelihood of the swaths  $\mathbf{X}$  and their target label masks  $\mathbf{y}$ , rewritten as:

$$p_\theta(\mathbf{y}, \mathbf{X}) = p(\mathbf{y} | \mathbf{X})p(\mathbf{X}) \quad (4)$$

$$= p(\mathbf{y} | \mathbf{z})p(\mathbf{z}) \left| \det \left( \frac{d\mathbf{z}}{d\mathbf{X}} \right) \right| \quad (5)$$

$$= p(\mathbf{y} | \mathbf{z})p(\mathbf{z}) \prod_{i=1}^D \left| \det \left( \frac{d\mathbf{h}_i}{d\mathbf{h}_{i-1}} \right) \right| \quad (6)$$

where the last two terms are given by the change of variable formula. As supervision is not provided for all the tiles,  $p(\mathbf{y} | \mathbf{z})$  in Eq. (6) can be estimated only for tiles from the labeled set  $\mathcal{L} = \{(x_k, y_k)\}_{k=1}^K$ . For the tiles of the unlabeled set  $\mathcal{U} = \{x_k\}_{k=1}^K$ , the label entropy is minimised instead, to promote sharp predictions. The overall objective function takes the following form, by equivalently maximizing the log likelihood:

$$\begin{aligned} \max \quad & \sum_{(x_k, y_k) \in \mathcal{L}} (\log p(y_k | z_k) + \log p(z_k)) + \sum_{x_k \in \mathcal{U}} \left( \sum_{y \in \mathcal{Y}} \log p(y | z_k) + \log p(z_k) \right) + \\ & + \sum_{j=1}^D \log \left| \det \left( \frac{dh_j}{dh_{j-1}} \right) \right| \end{aligned} \quad (7)$$

<sup>5</sup>label entropy is minimised instead of the cross-entropy whenever supervision is not available.

Table 3: Machine-Learning-based validation: IResNet classification results on test set.

	Ci	As	Ac	St	Sc	Cu	Ns	Dc	Mean
Accuracy (%)	81.30	84.50	88.29	97.73	88.90	92.40	90.92	98.84	90.36
F1 score	0.68	0.43	0.45	0.58	0.80	0.40	0.47	0.58	0.55
IoU index	0.27	0.20	0.17	0.15	0.50	0.06	0.18	0.18	0.21

Table 4: Physical-based validation: IResNet classification results for liquid water path (LWP) and cloud optical thickness (COT) distributions. We compare the distributions predicted by the IResNet with the ground truths using both Kullback-Leibler (KL) divergence and the Wasserstein (W) distance.

	metric	Ci	As	Ac	St	Sc	Cu	Ns	Dc
LWP	KL div.	$0.11 \cdot 10^{-1}$	$0.38 \cdot 10^{-2}$	$0.71 \cdot 10^{-2}$	$0.32 \cdot 10^{-1}$	$0.33 \cdot 10^{-2}$	$0.39 \cdot 10^{-1}$	$0.23 \cdot 10^{-1}$	$0.35 \cdot 10^{-1}$
	W dist.	$0.12 \cdot 10^{-3}$	$0.54 \cdot 10^{-4}$	$0.76 \cdot 10^{-4}$	$0.18 \cdot 10^{-3}$	$0.51 \cdot 10^{-4}$	$0.21 \cdot 10^{-3}$	$0.95 \cdot 10^{-4}$	$0.15 \cdot 10^{-3}$
COT	KL div.	$0.47 \cdot 10^{-2}$	$0.14 \cdot 10^{-1}$	$0.34 \cdot 10^{-1}$	$0.34 \cdot 10^{-1}$	$0.13 \cdot 10^{-1}$	$0.28 \cdot 10^{-1}$	$0.57 \cdot 10^{-1}$	$0.47 \cdot 10^{-1}$
	W dist.	$0.46 \cdot 10^{-4}$	$0.58 \cdot 10^{-4}$	$0.68 \cdot 10^{-4}$	$0.21 \cdot 10^{-3}$	$0.38 \cdot 10^{-4}$	$0.99 \cdot 10^{-4}$	$0.17 \cdot 10^{-3}$	$0.73 \cdot 10^{-4}$

### 3.3 Results

We randomly split the labelled tiles into training (70%), validation (10%) and test (20%) sets. We report test classification accuracies, F1 score and Intersection over Union index, per class and on average, in Table 3 for the model with the best mean accuracy on the validation set. Figure 2 shows the predictions obtained over one day of images and the occurrences (gridded at  $0.5^\circ$  latitude  $\times$   $0.5^\circ$  longitude) of three predicted classes over the month. Predicted classes (Fig. 2a) are spatially contiguous across swaths (this is not a constraint of our algorithm). The occurrences appear spatially coherent with Sc clouds occurring mostly over upwelling regions of the major oceans (Fig. 2c); Dc clouds are more confined to equatorial regions (Fig. 2d); and Ci (high) clouds more globally widespread (Fig. 2b). The highest occurrence of Ci clouds appears roughly over the inter-tropical convergence zone and is spatially correlated with Dc clouds, in agreement with Mace et al. [2006]. Most interestingly, the heatmaps show great spatial similarities with the ones reported by Sassen et al. [2008a,b, 2009], where the authors studied occurrences of cloud classes labeled by CloudSat over a period of one Sassen et al. [2008a,b] and two years Sassen et al. [2009]. All occurrences are shown in Fig. 3.

As an additional physical-based evaluation we consider the distributions of the liquid water path (LWP) and cloud optical thickness (COT) variables for the predicted classes and the ground truth given by CloudSat<sup>6</sup>. LWP defines the total amount of liquid water present in the whole atmospheric column on a given point. COT is a measure of the thickness between the bottom and top of a cloud. Qualitatively, differences between the predicted distributions and the ground truth are minimal for both variables. Results are shown in Figure 4. In Table 4 we also provide a quantitative comparison between predicted distributions and CloudSat ones by means of both the Kullback-Leibler divergence and the Wasserstein distance. We note that the class "Deep Convection" is the one associated with largest accuracy (98.84%) despite having relatively large differences with the ground truth values in both the LWP and COT variables (see Table 4).

In general, as the available supervision is minimal, we argue that it is inadequate to gauge the quality of a method considering exclusively its accuracy on testing samples. Therefore, a physical-based evaluation is more tailored to cloud classification studies, it does not suffer from the minimal supervision and it should always complement the more basic metrics discussed in Table 3.

### 3.4 Conclusions

In this work we first proposed CUMULO, a new benchmark dataset for training and evaluating global cloud classification models. It consists of one year of 1km resolution MODIS hyperspectral images merged with pixel-width tracks of CloudSat labels. We think that this is an important step for engaging the Machine-Learning community to develop innovative methods and solutions to climate related problems. In particular,

<sup>6</sup>Note that both LWP and COT features were not used for training.

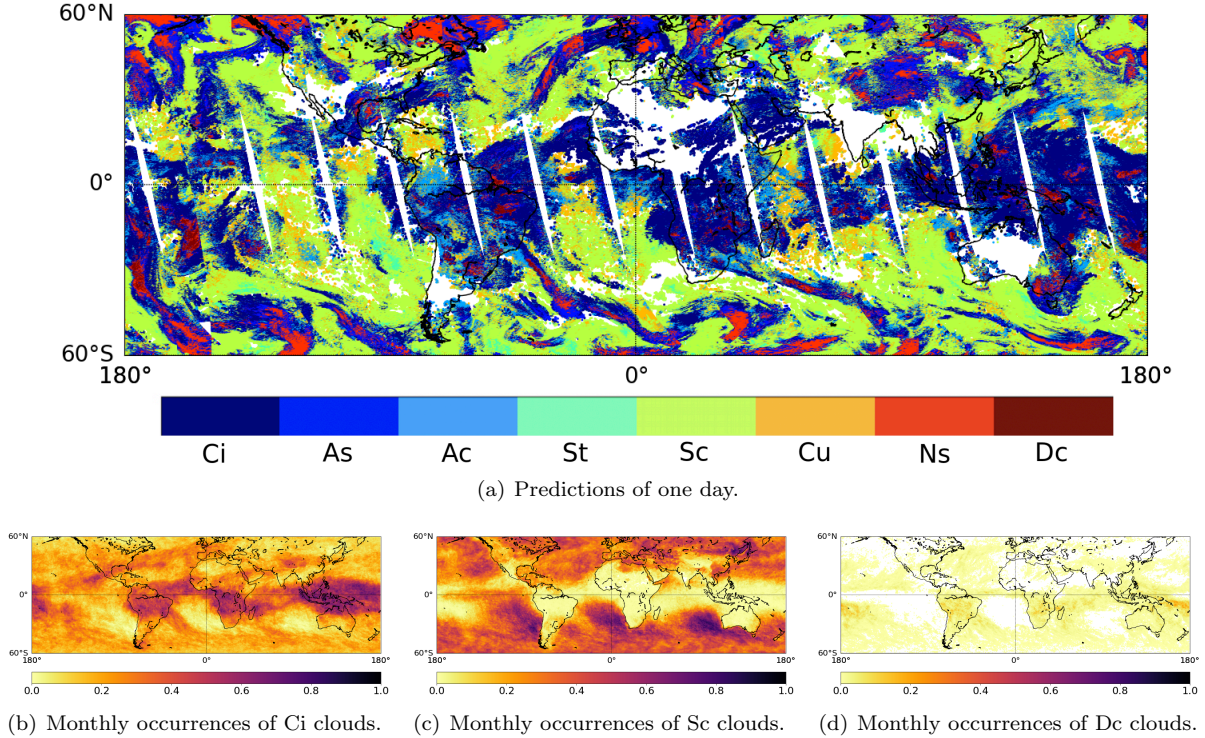


Figure 2: IResNet classification results. Occurrences are computed for January 2008. Occurrences of all classes are shown in Figure 3.

the proposed dataset presents several important challenges: (a) labels make up for less than 1% of the data (weakly-labelled data), (b) a pixel can have multiple labels (multi-labelled data), (c) certain type of clouds (i.e., deep convection) are underrepresented (class imbalance). Moreover, within a single cloud class we can still distinguish a rich variety of cloud organizations at the mesoscale (from 5 to several hundred kilometers) and slightly different physical properties, implying the existence of sub-classes for each given class. Proposing novel unsupervised models that directly discover fine-grained classes with only access to the observed coarse labels could be an important new line of research for both the Climate and Machine-Learning community.

We also provided a first high resolution spatiotemporal cloud classification baseline performance on CUMULO. To complement the standard ML prediction scores, we made use of the validation channels of CUMULO to analyse the results from a physical perspective. First, the occurrences over the analyzed month are found to be qualitative similar to previous studies. Second, we found that our baseline results are physically reasonable in terms of the liquid water path and cloud optical thickness distributions of the predicted classes. The reported analysis is quantitative and physical, but limited to January 2008. Since CloudSat needs 16 days to complete a cycle, we leave a rigorous comparison of the (predicted) monthly occurrences for future studies, using 1 year of classification.



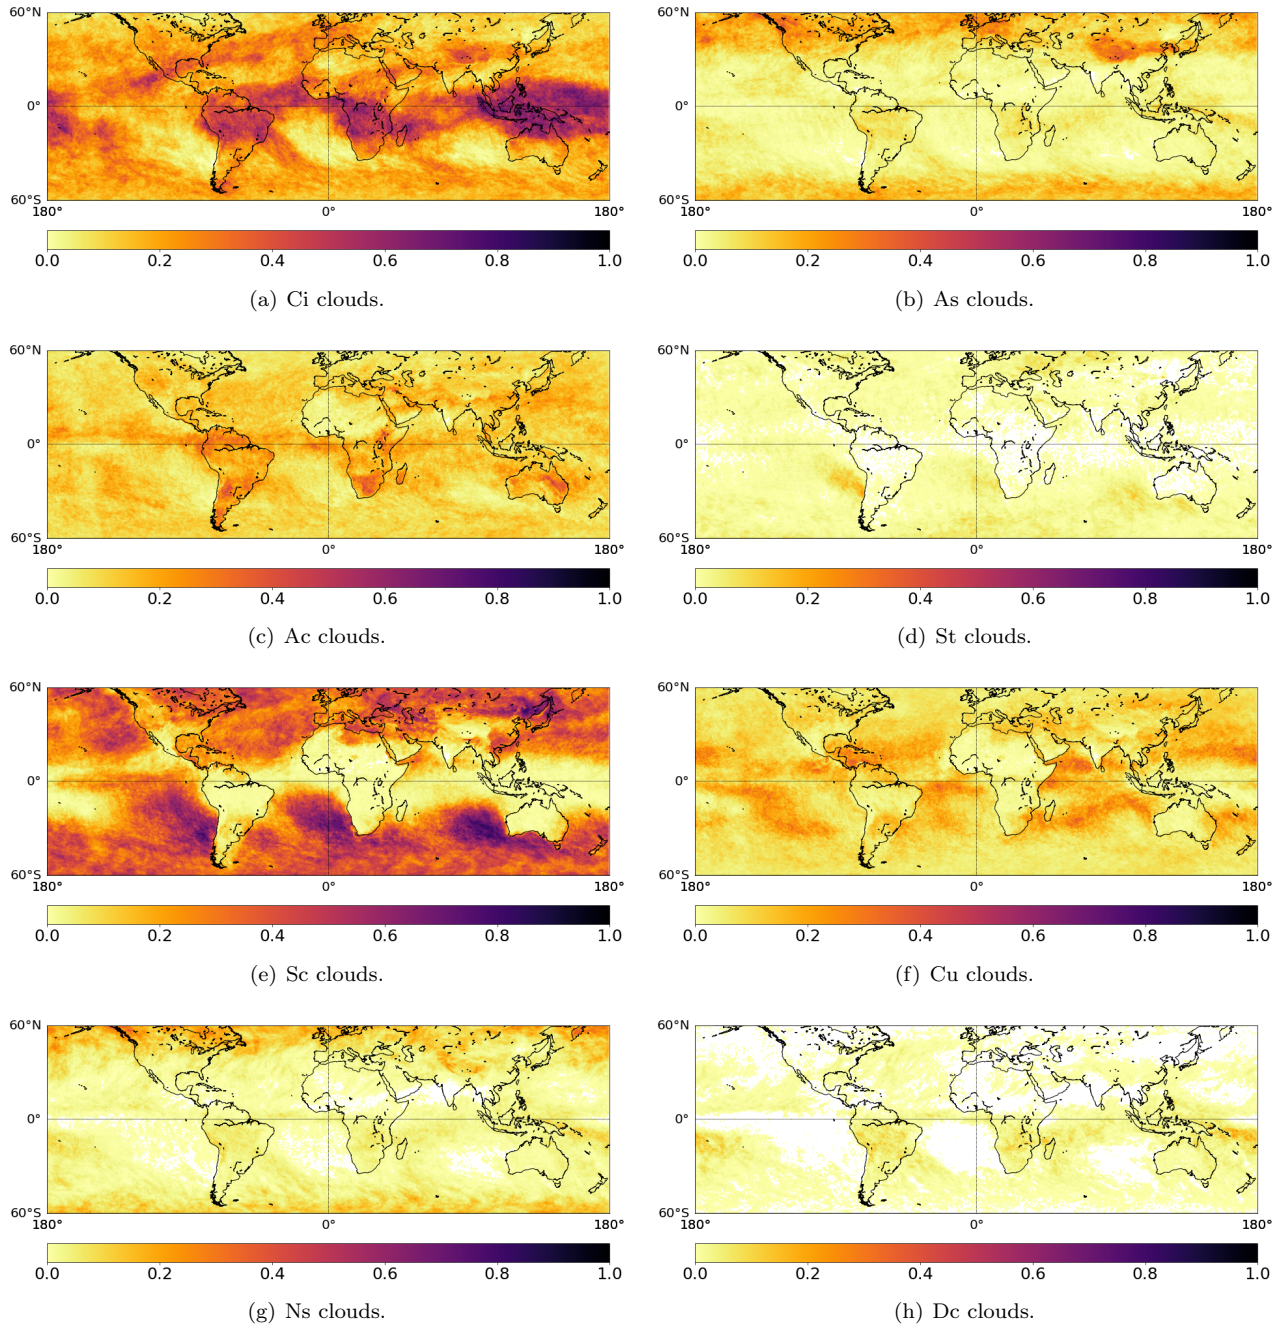


Figure 3: Occurrences of the cloud classes predicted by IResNet for January 2008.



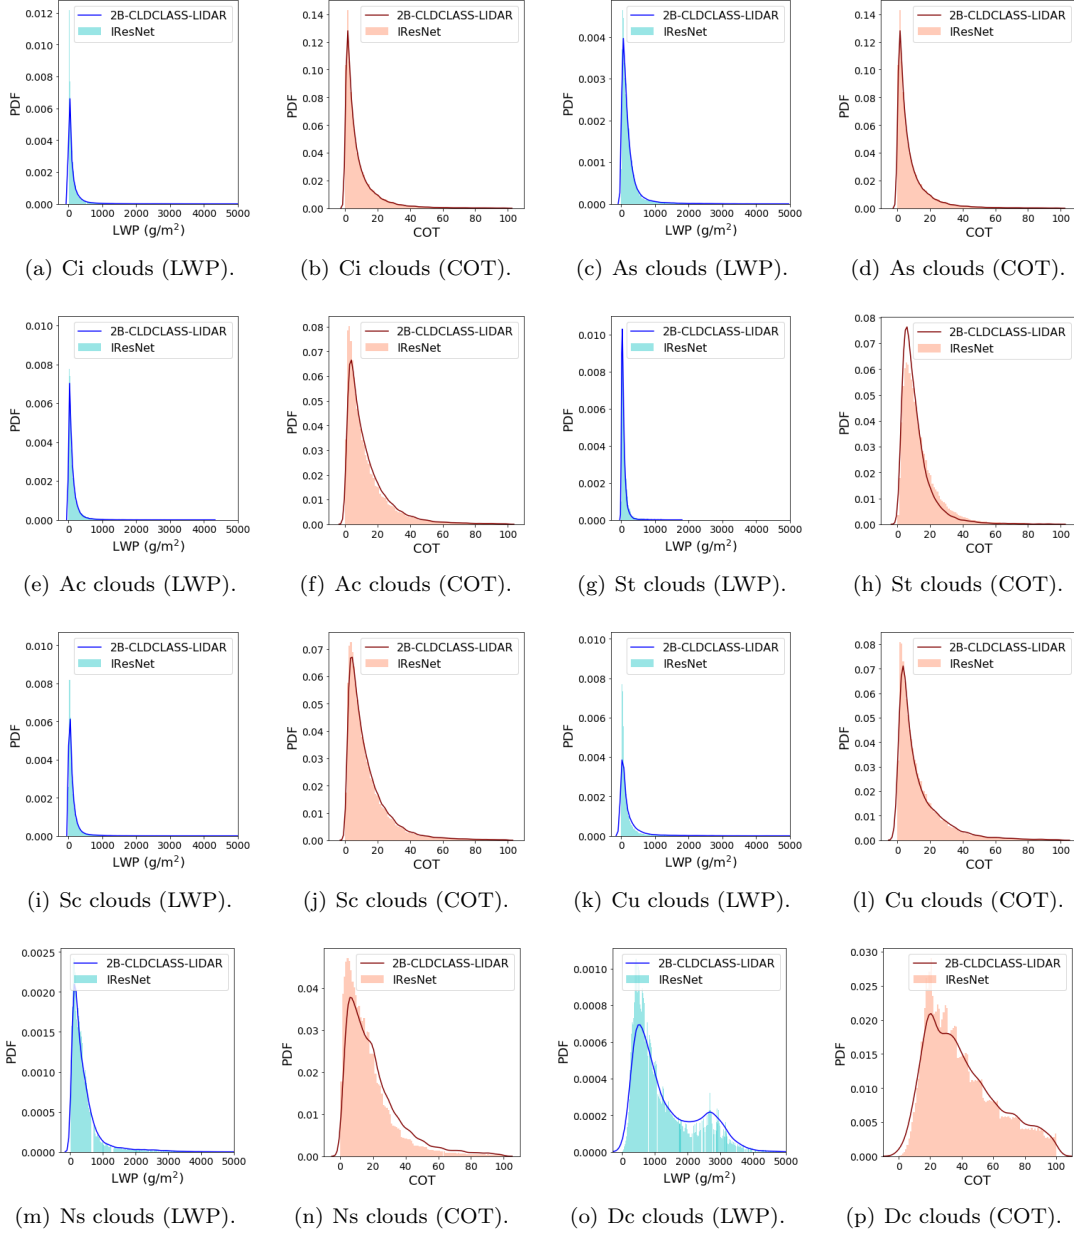


Figure 4: PDFs of the liquid water path (LWP) and cloud optical thickness (COT) for the cloud classes predicted by the IResNet for January 2008. LWP(COT) histograms are drawn in blue(red).

## Data availability

The code used for extracting CUMULO is hosted at <https://github.com/FrontierDevelopmentLab/CUMULO>. The dataset will soon be made publicly available. In the meantime, please contact us for any request.

## Acknowledgments

This work is the result of the 2019 ESA Frontier Development Lab <sup>7</sup> (FDL) Atmospheric Phenomena and Climate Variability challenge. We are grateful to all organisers, mentors and sponsors for providing us this opportunity. We thank Google Cloud for providing computing and storage resources to complete this work. Finally, we thank Yarin Gal for helpful discussions and Sylvester Kaczmarek for his help and support in coordinating the work.

## References

- Reno Knutti, Maria A. A. Rugenstein, and Gabriele C. Hegerl. Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10:727–736, 2017.
- Leon D. Rotstayn and Mark A. Collier. Why does aerosol forcing control historical global-mean surface temperature change in cmip5 models? *Journal Of Climate*, 28:66086625, 2015.
- Graeme L Stephens. Cloud feedbacks in the climate system: A critical review. *Journal of climate*, 18(2): 237–273, 2005.
- Sandrine Bony, Bjorn Stevens, Dargan MW Frierson, Christian Jakob, Masa Kageyama, Robert Pincus, Theodore G Shepherd, Steven C Sherwood, A Pier Siebesma, Adam H Sobel, et al. Clouds, circulation and climate sensitivity. *Nature Geoscience*, 8(4):261, 2015.
- Tapio Schneider, Colleen M Kaul, and Kyle G Pressel. Possible climate transitions from breakup of stratocumulus decks under greenhouse warming. *Nature Geoscience*, 12(3):163, 2019.
- William B. Rossow and Robert A. Schiffer. Isccp cloud data products. *Bulletin of the American Meteorological Society*, 72(1):2–20, 1991. doi: 10.1175/1520-0477(1991)072<0002:ICDP>2.0.CO;2. URL [https://doi.org/10.1175/1520-0477\(1991\)072<0002:ICDP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1991)072<0002:ICDP>2.0.CO;2).
- Andreas Muhlbauer, Isabel L McCoy, and Robert Wood. Climatology of stratocumulus cloud morphologies: microphysical properties and radiative effects. *Atmospheric Chemistry and Physics*, 14(13):6695–6716, 2014.
- Chengwei Zhang, Xiaoyong Zhuge, and Fan Yu. Development of a high spatiotemporal resolution cloud-type classification approach using himawari-8 and cloudsat. *International Journal of Remote Sensing*, 40(16): 6464–6481, 2019.
- Stephan Rasp, Hauke Schulz, Sandrine Bony, and Bjorn Stevens. Combining crowd-sourcing and deep learning to understand meso-scale organization of shallow convection. *arXiv preprint arXiv:1906.01906*, 2019.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- 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*, pages 234–241. Springer, 2015.

---

<sup>7</sup><https://fdleurope.org/>

- G Papandreou, L-Ch Chen, K Murphy, and AL Yuille. Weakly-and semi-supervised learning of a dcnn for semantic image segmentation (2015). *arXiv preprint arXiv:1502.02734*.
- Qizhu Li, Anurag Arnab, and Philip HS Torr. Weakly-and semi-supervised panoptic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 102–118, 2018.
- Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Hybrid models with deep and invertible features. *arXiv preprint arXiv:1902.02767*, 2019.
- Gerald Mace, Min Deng, Brian Soden, and Ed Zipser. Association of tropical cirrus in the 1015-km layer with deep convective sources: An observational study combining millimeter radar data and satellite-derived trajectories. *J. Atmos. Sci.*, 63:480503, 2006.
- Kenneth Sassen, Zhien Wang, and Dong Liu. Global distribution of cirrus clouds from cloudsat/cloud-aerosol lidar andinfrared pathfinder satellite observations (calipso) measurements. *Journal of Geophysical Research: Atmospheres*, 113(D00A12), 2008a.
- Kenneth Sassen, Zhien Wang, and Dong Liu. Classifying clouds around the globe with the cloudsat radar: 1-year of results. *Journal of Geophysical Research: Atmospheres*, 113(D8), 2008b.
- Kenneth Sassen, Zhien Wang, and Dong Liu. Cirrus clouds and deep convection in the tropics: Insights from calipso and cloudsat. *Journal of Geophysical Research: Atmospheres*, 114(D00H06), 2009.