

Supporting information for

Capturing the diversity of mesoscale trade wind cumuli using complementary approaches from self-supervision

Dwaipayan Chatterjee¹, Sabrina Schnitt¹, Paula Bigalke¹, Claudia Acquistapace¹, Susanne Crewell¹

¹Institute for Geophysics and Meteorology, University of Cologne, Cologne, Germany

February 16, 2024, 12:06pm

Contents of this file

1. S1 Domain description
2. S2 Network architectures
3. S3 Determination of optimal cluster number
4. S4 Sensitivity tests
5. S5 Attention maps of N1
6. S6 Environmental characteristics
7. S7 Transition probability

Corresponding author: D. Chatterjee, Institute for Geophysics and Meteorology, University of Cologne, Cologne, Germany. (dchatter@uni-koeln.de)

S1 Domain description

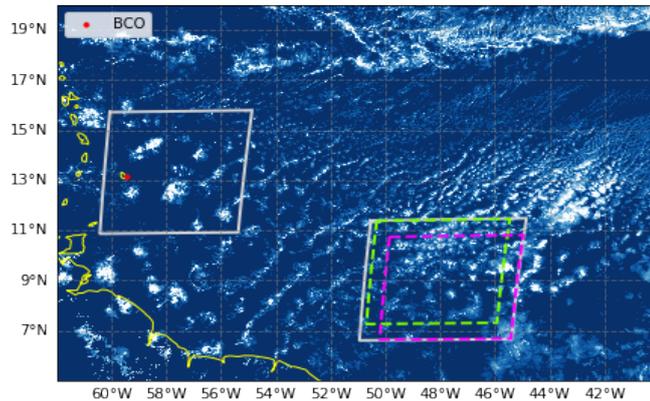


Figure S1 (Domain). GOES's COD image on February 2, 2020, at 13:00 UTC with coastal boundaries (thick yellow) and Barbados Cloud Observatory (red dot). One (out of five) random and a fixed (Barbados domain) 256 x 256-pixel crop over EUREC⁴A domain are shown. During the learning process, each crop is twice randomly sub-cropped (pink and green dashed lines) by the network, leading to a spatial dimension of 75% (192 x 192 pixels) of the original crop. The Barbados domain enables comparison with ground-based measurements in future studies.

S2 Network architectures

1. Continuous network (N1)

1.1 Definition of the network input

N satellite images of COD built the input training data set $X = \{x_1, x_2, x_3, \dots, x_N\}$ of the deep learning architecture illustrated in Fig. S2.1 (Schematic diagram of N1). The only intuitive augmentation we opt for here is global random cropping for learning continuous representations. For random cropping, we opt for two global crops (x_1, x_2) with a random 0.75 fraction (192 x 192 pixels) of the parent satellite image to focus on the global distribution of the cloud system. Figure S2.1 (Schematic diagram of N1) shows two random crops (teacher t and student s) are fed into different branches of the network. Therefore, it becomes challenging for one side of the network to know what part of the parent satellite image the other is being fed with; therefore, during the learning process, it focuses on the critical semantics of global cloud distribution.

1.2 General network architecture

The neural network's task is to learn visual features from each satellite image. A function g represents the transformations performed by the network's vision Transformer (ViT) as $g(x_i) = h_j$ with $i = 1 \dots N, j = 1 \dots M$ that maps the image x_i into the array of features $h = \{h_1, h_2, h_3, \dots, h_M\}$, where M is the output dimension of ViT feature arrays. The selected dimension of M is equal to 384, which means the information contained in the 192 x 192 satellite observation space is being non-linearly dimensionally reduced to 384 vector space. ViT is a sequence of self-attention (Vaswani et al., 2023), and feed-forward layers paralleled with skip connections. The mechanism of ViT (Dosovitskiy et al., 2021) takes non-overlapping contiguous image patches of resolution $N \times N$ pixels, where $N=16$ for this work, along with their positional encoding as an input. Without

the positional encoding, the output feature vector from ViT is invariant to the arrangement of these $N \times N$ patches. Meanwhile, with positional encoding, it learns the relative position of the objects in the image. Therefore, the model learns the relationship between the patches, and thus,

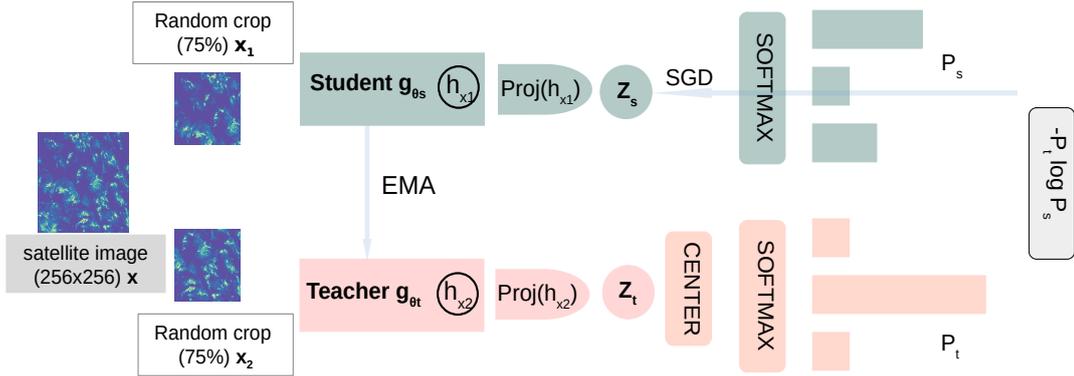


Figure S2.1 (Schematic diagram of N1). This work adopts a deep learning architecture from Caron et al. (2021), where x_1 and x_2 are 75% random crops of the parent satellite image x . The student and teacher vision transformers ($g_{\theta_{s/t}}$) have the same number of trainable parameters (weights and biases) θ . The feature output h_{x_i} from g_{θ} subsequently connects to $Proj(h_{x_i})$, a 3-layer multilayer perceptron. Softmax (Bridle, 1989) normalizes MLP's raw activation ($Z_{s/t}$), and centering maintains teacher activations (Z_t) near batch mean properties. P_s and P_t represent normalized student distribution of Z_s and centered and normalized distribution of teacher activation Z_t . Back propagation in student network optimizes its parameters through stochastic gradient descent (SGD), minimizing cross-entropy between P_t and P_s . Teacher parameters (g_{θ_t}) are exponential moving averages (EMA) of students (g_{θ_s}), aligning the networks.

it learns how a particular arrangement of cloud distribution usually occurs. The ViT architecture can identify long-range spatial dependencies (Khan et al., 2022) by learning relevant information in the image. The activation function used in the ViT is Gaussian error linear units (Hendrycks & Gimpel, 2023) (GELU), as the GELU function behaves smoother than other activation functions when values are closer to zero and thus is more effective at learning complex patterns in the data.

Further, h_j is non-linearly projected to Z_l with $l = 1 \dots L$ using a three-layer multilayer perception (MLP) (Rumelhart et al., 1986) activated by GELU followed by l_2 normalization and a linear layer. Here $Z = \{Z_1, Z_2, Z_3, \dots, Z_L\}$ is the final output dimension of the pipeline. The feature space dimensions are decided based on input dimensions, the complexity of information context, and neural network complexity. Caron et al. (2021) suggest that if the training dataset size is much less than 1.3 million, then the final dimensions of Z_l should be reduced compared to the default dimensions $L = 65536$. We tested rather two different values ($L = 128$ and $L = 8192$) and found a much better suitability of the larger value from visual inspection. Our aim here is not to find the optimal feature vector size but a functional size that can optimize the network and smoothly converge the training. Therefore, the optimal dimension size of the dimensionally reduced images in self-supervised learning is not the focus of this work. Figure S2.1 (Schematic diagram of N1) shows two different branches in the network: student and teacher. The point to note here is that they have the same general architecture and pipeline, but the parameters (weights and biases) learned during training are different.

1.3 Upper branch of the network (Student)

The upper branch of the network, represented in Figure S2.1 (Schematic diagram of N1), by the student transformer g_s and further projected by MLP, ingests one random augmented global

crop of the parent satellite image and outputs feature vector Z_s . Softmax normalizes Z_s such that all values are between 0 and 1, and the integration over all $L = 8192$ elements of its probability distribution yields 1. This probability distribution of the feature vector Z_s is input to the cross entropy loss function described later. The soft-max probability for an input x_i of the student network can be described as

$$p_s^{(i)} = \frac{\exp(\frac{1}{\zeta_s} Z_s^{(i)})}{\sum_{m=0}^k \exp(\frac{1}{\zeta_s} Z_s^{(m)})} \quad (1)$$

where ζ_s is the temperature parameter for the student network and is set to 0.1. The ζ parameter controls the sharpening of the probability distribution. A higher value of ζ implies smoothed probability.

1.4 Lower branch of the network (Teacher)

The lower branch of the network represented in Figure S2.1 (Schematic diagram of N1) by the teacher transformer applies function g_t to the other remaining global crop of the parent satellite image, and the MLP projects outputs feature vector Z_t . Unlike P_s , before normalizing P_t individually with soft-max, vector Z_t is centered around the mean properties of all images in a batch. A batch refers to the number of samples propagating through the neural network before updating the model parameters. Centering is done to prevent any feature from dominating, as the mean will be somewhere in the middle of the batch sample properties. While applying the temperature ζ_t parameter for the teacher, it is kept lower ($\zeta_t = 0.05$) to sharpen the probability of Z_t artificially. Therefore, the feature vector Z_t of the teacher branch is centered and sharpened before it becomes input for the loss function.

1.5 Cross entropy loss of the network

When the feature vectors of the two branches capture similar information from the global crops of the satellite parent image, the loss becomes lower and vice-versa. That’s how the network branches are encouraged to focus on the common image characteristics, progressively making the feature vectors similar.

$$\min_{\theta_s} \sum_{x \in (x_1, x_2)} P_t(x) \log(P_s(x)) \quad (2)$$

This is achieved through the cross-entropy loss function applied on the centered and sharpened probability distribution of the teacher branch P_t and smoothed distribution of the student branch P_s . As shown in equation 2, the loss function minimizes θ_s , i.e., the student network’s parameters (weights and biases). Teacher network parameters or P_t guide the student network during the training phase, as discussed in Subsection 1.6.

1.6 Optimization for convergence

The loss function minimization happens progressively layer by layer, derivating the loss function with respect to θ_s parameters and adjusting parameter values in each layer by backpropagation. At the end of the minimization, we obtain a configuration of parameters for the student network that will be ready for the next iteration with a new batch of images. Stochastic gradient descent (Bottou, 2012) (SGD) is only applied to the student network parameters θ_s , and the teacher parameters θ_t are built through past iterations of the student network (Caron et al., 2021). As shown in equation 3, θ_t is the exponential moving average (EMA) of θ_s with λ following a cosine scheduled from 0.996 to 1 during training.

$$\theta_t = \lambda \theta_t + (1 - \lambda) \theta_s \quad (3)$$

During optimization, a collapse can occur regardless of the input provided to the model; the output becomes constant or is predominantly influenced by a single dimension. In other words, the model's predictions across different dimensions or features become uniform, leading to zero ideal loss value. Therefore, centering and sharpening introduced in Subsection 1.3 and 1.4 and EMA (Subsection 1.6) are the easiest acceptable ways to prevent collapsing in the described teacher-student framework.

1.7 Training and libraries

To set up this architecture, we use the software package DINO from Facebook Artificial Intelligence Research (FAIR) (Caron et al., 2021) based on PyTorch. The open-source VISSL computer vision library (Goyal et al., 2021) adapted the DINO neural network to our requirements. Based on sensitivity tests on training loss, visualization of dimensionally reduced feature space, and ablation study of the original network on longer training showing improving performance, we train the model up to 800 epochs. Training the neural network for 800 epochs on 4 V100 GPUs took 16.5 hours or 66 core hours.

2 Discrete network (N2)

We briefly describe the functional mechanism of the discrete neural network (N2) and its learning scheme. Refer to Section 3 from (Chatterjee et al., 2023) for a detailed network description. The data loading nature of N2 remains the same as of N1 (Subsection 1.1 of S2). The general architecture has a pipeline similar to the continuous approach set up, with the image processing backbone here being a convolutional residual network with 50 layers of depth (ResNet-50, (He et al., 2015)), followed by a projection head of MLP with ReLU activations (Fukushima, 1975) and a linear layer.

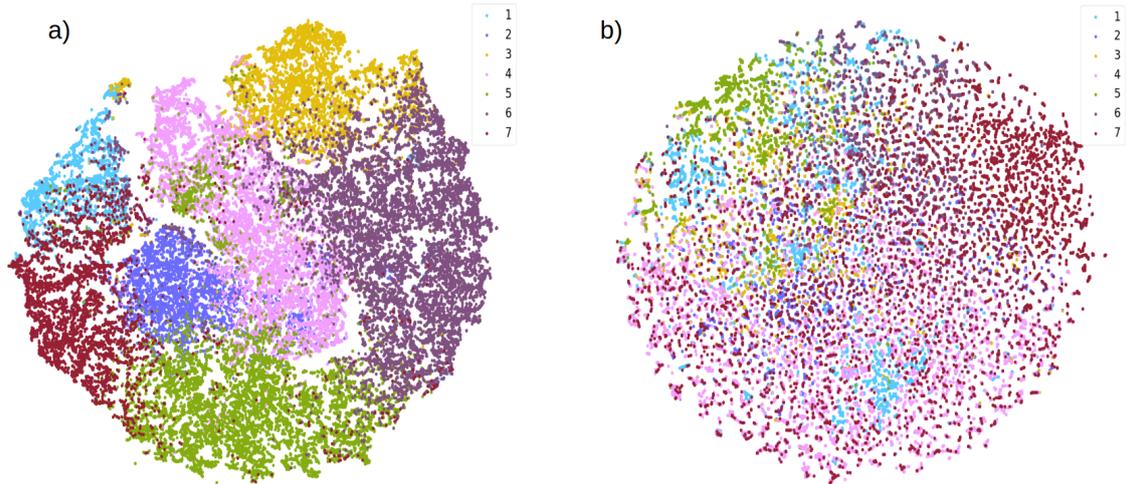


Figure S2.2 (N2 outputs). a) Sparse 2D feature space obtained from N2 by applying the tSNE algorithm on z_x features of 51,000 satellite images. The perplexity and epsilon derived from auto-configuration for t-SNE runs are 30 and 1150, respectively. b) Same as Figure 1.b in the main text, this uses direct clustering on satellite images with N2, overlaying labels on the continuous feature space from N1 for comparison.

Therefore, similar to Figure S2.1 (Schematic diagram of N1), for the upper branch, the features obtained at the end of the pipeline (like Z_s in the continuous approach) are clustered using spherical k-means (where $k=7$), and features are allocated a pseudo-label (T) according to their closest centroid. Further, the features obtained from the lower branch are compared with the calculated upper branch centroids using cosine distance (D_T). Finally, T from the upper branch and D_T from the lower branch are inputs of the cross-entropy loss function as discussed in Subsection 1.5 of S2 and are progressively minimized during training. We call the labels pseudo-labels during the training stage as they can change to minimize the loss function better. Finally, at the end of the training, we collect the labels for each satellite image and further evaluate their separation using auxiliary datasets.

S3 Determination of optimal cluster number

We apply the following metrics to two-dimensionally reduced representations (using tSNE) on h_j from N1 to identify the best optimal cluster:

Distortion metric: The distortion metric considers the cluster's tightness by computing the sum of squared distances (SSD) from each point to its assigned center, which tends to decrease toward 0 as we increase the number of clusters (K). This shows an exponential shape leveling off such that the shape of the curve results in an elbow, but the optimal cluster or the point of inflection represents the point where adding additional clusters stops adding useful information. Also, adding clusters beyond the inflection point also makes the clusters harder to separate; thus, we start to observe diminishing returns by increasing k. The elbow blue line curve in figure S3 (Metric scores) shows $k = 7$ as the sweet spot of optimal clustering.

Silhouette metric: Apart from taking cluster closeness into account, this metric also considers distances between points of one cluster and the nearest other cluster center. This means that in order to have a good silhouette score, clusters generally need to be tighter and farther apart from each other. If the Silhouette coefficient for each point is close to 0, it means that the point is between two clusters; if it is close to -1, then that point is in the wrong cluster, and if it is close to +1, it is in the correct cluster. The average silhouette coefficient calculated for all 51,000 samples shows two local maxima at values of 0.37 ($k=3$) and 0.36 ($k=7$), as shown in Figure S3 (Metric scores). Note that the values are not close to one, meaning the cluster doesn't lie very far from each other, further suggesting the continuous nature of cloud organizations.

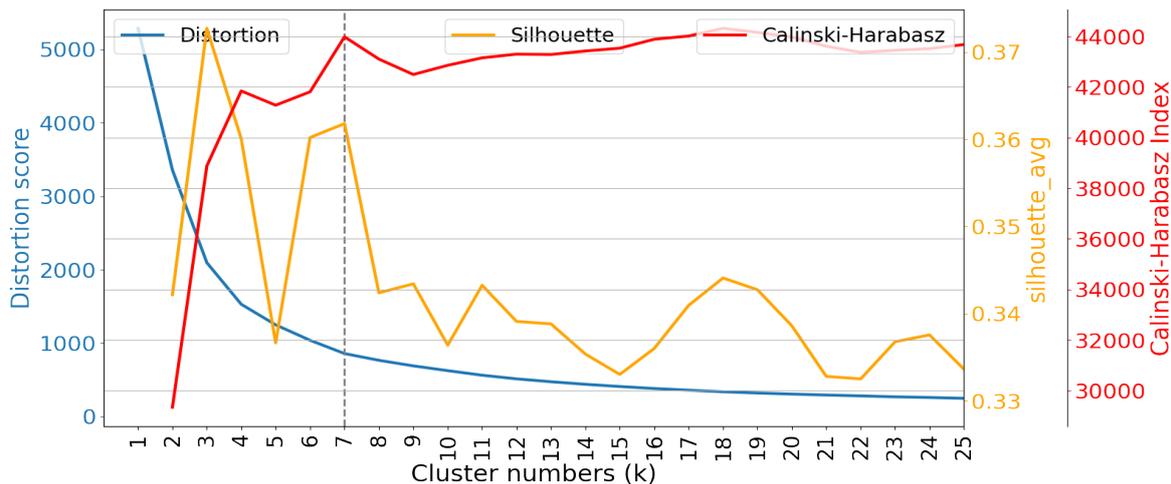


Figure S3 (Metric scores). Results of three different metric scores of distortion, silhouette, and Calinski-Harabasz, shown along with varying cluster numbers along the abscissa. The vertical-dashed line is drawn at cluster 7, which shows the chosen inflection point for the optimal cluster.

Calinski-Harabasz metric: In comparison, the Calinski-Harabasz metric assesses the separation and compactness of the clusters. It denotes the ratio of the sum of inter-cluster dispersion and the sum of intra-cluster dispersion for all clusters. A good clustering result has a high Calinski-Harabasz Index value. The maximum lies at cluster 7, having a score of 43000.

In summary, the two metrics direct towards $k=7$, and the difference between the two maxima ($k=3$ and 7) in silhouette is insignificant. Therefore, we take the common agreement of $k=7$ as the optimal cluster number and train N2 (Section 2 of S2) from scratch using 7 clusters.

S4 sensitivity tests

Here, we show that the choice of seven classes has passed several sensitivity tests, such as the dimensionality-reduction technique, size of the dataset, initial weights of the network, and different global crop sizes.

S4.1 Different data sample size

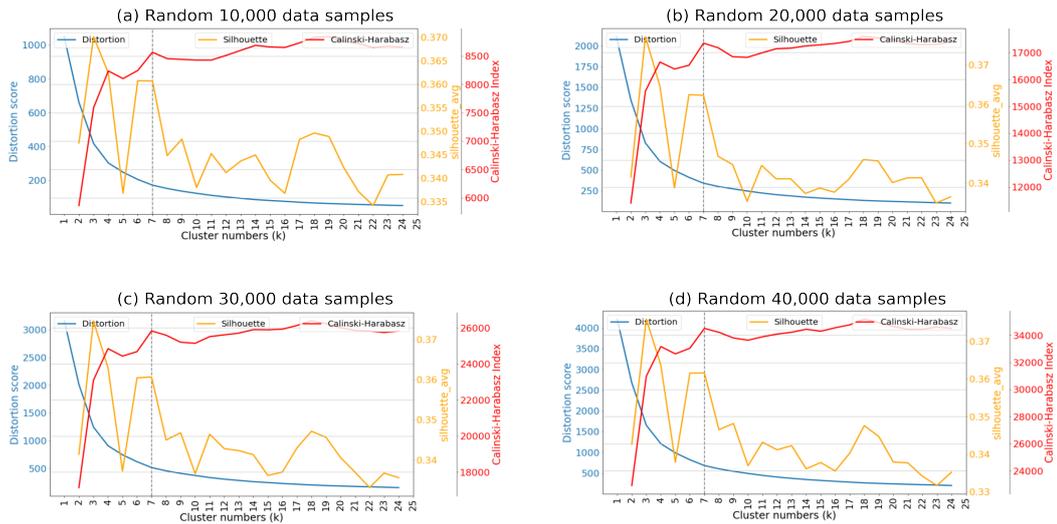


Figure S4.1 (Different data sample size). t-SNE initialized by PCA and using cosine distance as a unit of distance while constructing the two-dimensional space, on a sample size of a) 10,000, b) 20,000, c) 30,000 d) 40,000 data points.

Text S4.2 Dimensionality reduction techniques Manifold extraction algorithm is a generic term used for nonlinear dimensional reduction, or we can call them generalized PCA, which is sensitive to nonlinear structures in the data.

(Denby, 2020) train their neural network using ‘Euclidean distance’ as a metric in the loss function. However, for dimensionality reduction, use Isometric feature mapping (Isomap, (Tenenbaum et al., 2000)), an extension of kernel PCA. Isomap uses ‘geodesic distance’ as a measure of dis-

tance while reducing the dimensions. The geodesic distance looks for the shortest curve in the high dimensional space. Therefore, there are some limitations here:

1. Euclidean distance and high dimensions: First, we would like to point out that Euclidean distance breaks down in high dimensions (Aggarwal et al., 2001). Euclidean distance is sensitive to sparse data distribution in high dimensions. The direction becomes more critical since we normalize our feature vectors (the magnitude becomes one). Therefore, cosine distance (used in N1 and N2) is far more suitable for training.

2. Global versus local structures in the data: Second, we would like to mention that the Isomap algorithm considers maintaining the global pairwise distance (Gao et al., 2021). In other words, it neglects the local structure but only considers the global one.

Overcoming the limitations: t-SNE (van der Maaten & Hinton, 2008) preserves the neighboring local distances better.

1. We inject the global structure into our initialization of tSNE through PCA, which dictates which regions of the 2D space the points will appear.

2. Second, while reducing the dimensions, we keep cosine distance as a distance measurement criterion in t-SNE.

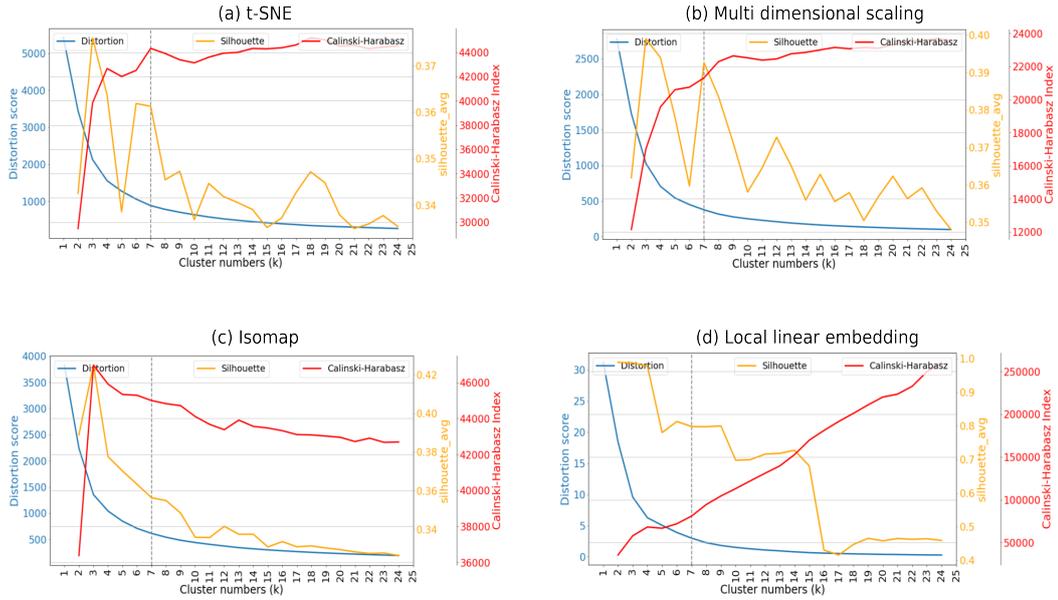


Figure S4.2 (Dimensionality reduction techniques). Different dimensionality reduction techniques applied on our high dimensional feature space for cluster number optimization, which is constructed using cosine distance (as unit of measuring distance between two feature vectors). See the main text for an explanation of the distortion, Silhouette, and Calinski-Harabasz algorithms. We can see the consensus between t-SNE and MDS, while Isomap and LLE could not come to a logical conclusion.

For our high dimensional feature space, we demonstrate optimal cluster number sensitivities to different dimensionality reduction techniques in Figure S4.2 (Dimensionality reduction techniques). In addition to t-SNE and Isomap, we also show multi-dimensional scaling (MDS) and local linear embedding (LLE) methods.

1. MDS, distance-wise, is more geometrically aware as it tries to preserve the inner products between the feature vectors.

2. LLE pays attention to only the local structure of the data. LLE also assumes that the high dimensional data is locally linear and a sample can be represented linearly by several samples in the neighborhood.

S4.3 Different initial weights of the network

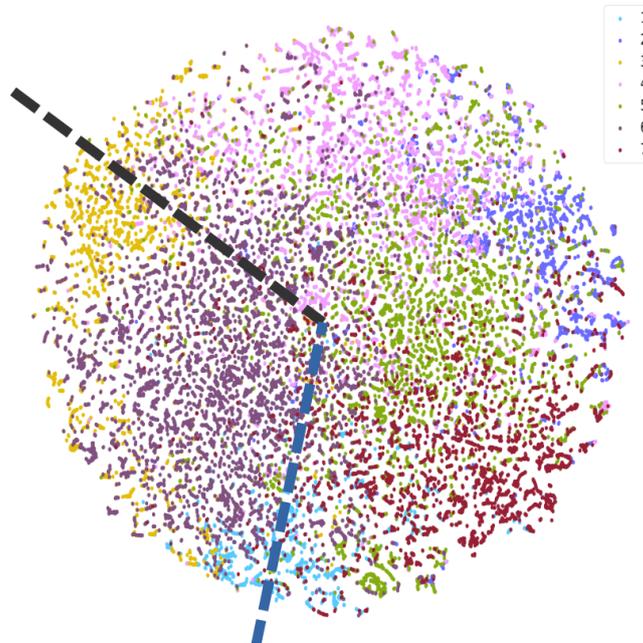


Figure S4.3 (Initial weights of the network). N1 and N2 initialized with different random seeds. As compared to main figure 1.b, we find only the orientation of the low dimensional projection to change where the black dashed line is the original location of class 1 and the blue dashed line is with a different random seed. The rest (sorting order and classification) remains the same.

S4.4 Different overlap threshold

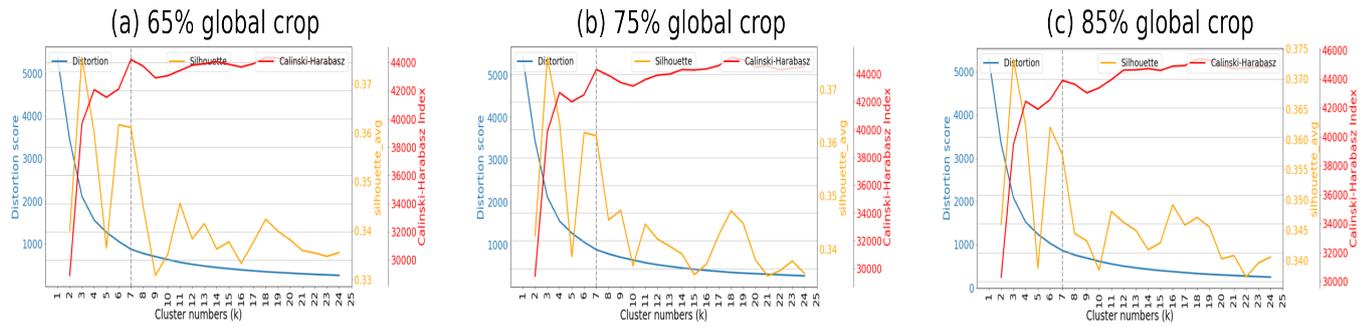


Figure S4.4 (overlap thresholds). Selecting different global crop size (65%, 75%, 85%) and training N1. t-SNE is initialized by PCA, and using cosine distance as a unit of distance, we find $k = 7$ as the common consensus among the three.

Therefore, Figure S4 (1-4) mainly argues about four things:

1. While reducing dimensions, it is important to consider how the original high-dimensional space was constructed. What distance was used while computing affinities? Then, use the same distance while constructing the lower dimensions.
2. It is important to consider a dimensionality reduction technique that considers both the local and global structure of the high dimensional data.
3. Varying number of samples from (10,000 to 50,000) still shows $k = 7$ as the most expressive optimal cluster number
4. Varying the Global crops from the 65th to 85th fraction shows crop size might not be a strong influencing factor in finding the optimal cluster number.

S5 Attention maps of N1

From a human perspective, cloud system distributions may appear to be relatively chaotic and noisy, and while trying to decide their visual characteristics, we may pay attention to some or all of the following: the organizational semantics of convective organization, the semantics of the clear sky regions, deep convective cell distributions, open and closed cells, and shallow convection distributions.

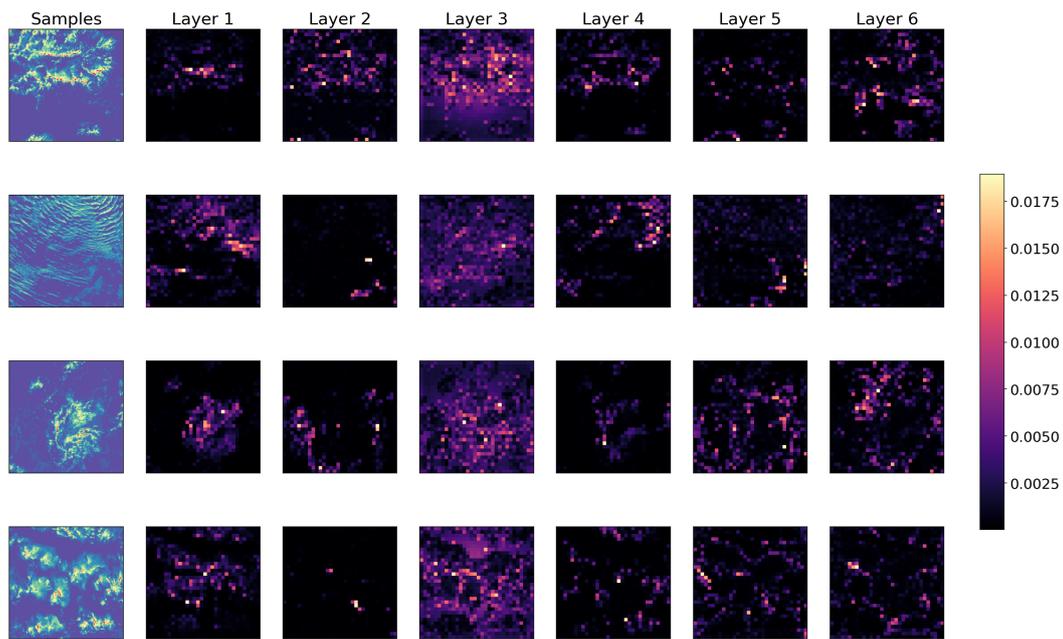


Figure S5 (Visualization of different layers). Four cloud systems with different organizations are selected as examples. Their respective self-attention maps from the final head of the teacher ViT (Figure S2.1 (Schematic diagram of N1)) show for the 6 layers of the self-attention head. The color bar indicates the range of the Gaussian error linear units (GELU) activation function for the activation maps. Higher values indicate more important features. All experiments are run with a default of six self-attention heads.

To better understand N1's decision and to build trust in the network's performance, it is crucial to see what the trained N1 architecture has learned to pay attention to when deciding the features of cloud system distributions.

Given a satellite image, the activation space in a neural network allows us to visualize whether a neuron should be activated, indicating what part of the image is important for the network. The self-attention layers in ViT try to decompose the input samples and learn relatively independent features. Thus, this experiment aims to see whether the activation space reveals the abstract patterns that we, as humans, can make sense of while deciding the feature's importance. In this setup, we use a single satellite image sample and pass it through the trained model, freezing the weights. The granularity ($N \times N$), or the number of pixels in a single patch, is controlled by the patch size, which is 8×8 pixels in this experiment. This is just for convenience and does not change the result compared to the 16×16 setting used for the main experiment.

Figure S5 (Visualization of different layers) shows that layer 1 activates at the dominant convective cells and deactivates at thin spread-out convection while layer 2 activates the thin spread convection. Layer 3 seems to try to learn and activate the clear sky features. In contrast to layer one, layer 4 activates the rest of the prominent convections. Like layer 2, layer 5 tries to look at the rest of the thin-spread convection. Layer 6 is uncertain and is not obvious to our eyes, and it may somehow try to deactivate for all the clear sky regions in the majority of cases and look for boundary semantics in the satellite image. Examining other example cases shows the same consistency, and therefore, it can be concluded that although the cloud system distributions are different, each attention map has learned to pay attention to relatively different, consistent,

sensible semantics of the cloud systems distribution and further indicates that we can trust the embedding space of the network.

S6 Environmental characteristics

Figure 2.c in Section 4.2 of the main manuscript showed the occurrences of 30 nearest neighbors of human-labeled satellite images (mentioned as human crops below) with machine-identified seven classes. Here, we aim to assess their existing environmental conditions.

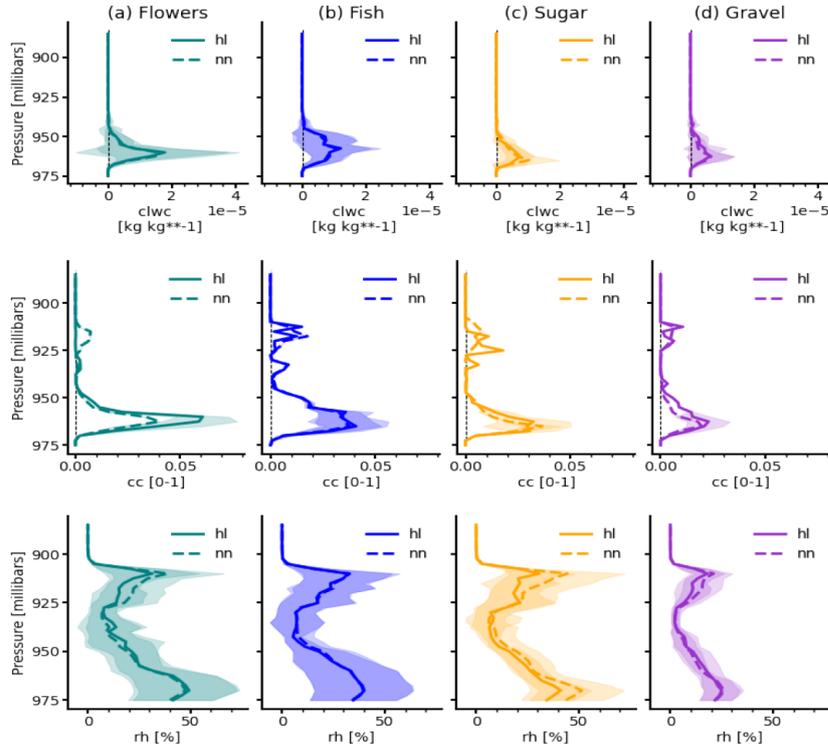


Figure S6. Comparison of 52 human labels (hl) environmental conditions with their nearest 30 neighbors (nn) using ERA-5. The top to bottom rows shows weighted-average and standard deviation profiles of cloud water content (clwc, kg kg^{-1}), cloud cover (cc), and relative humidity (rh, %) with the exception of cc variability shown in the interquartile range.

This complementary experiment can further help to trust human crops' relative positions in the feature space. If the human crops and the neighbors have a similar homogenous distribution of their physical properties, this implies that the human crops are in the consistent region of the feature space. Here, we take the ERA-5 vertical profile of cloud water content, cloud cover, and

relative humidity (Fig. S6) to compare the weighted averaged vertical profiles between human labels and their 30 nearest neighbors. When calculating these properties for human-labeled scenes, we weigh them with the level of agreement. In this way, the contribution of well-agreed organizations will contribute more than less agreed cloud organizations. We observe that there is hardly any difference in the vertical profiles except for the relative humidity of sugar and cloud cover for flowers. This may be due to quantitatively using 30 times more data.

S7 Transition probability

Based on our collected satellite imagery derived over Barbados cloud observatory (BCO, see Figure S1 (Domain)), we performed a transition probability analysis for all machine-identified cloud regimes. Here, we exploit the 10-minute temporal resolution of the data set to observe the transformation of cloud systems with time. Note that some of the temporal consecutive satellite imagery could be missing because of the pre-processing step.

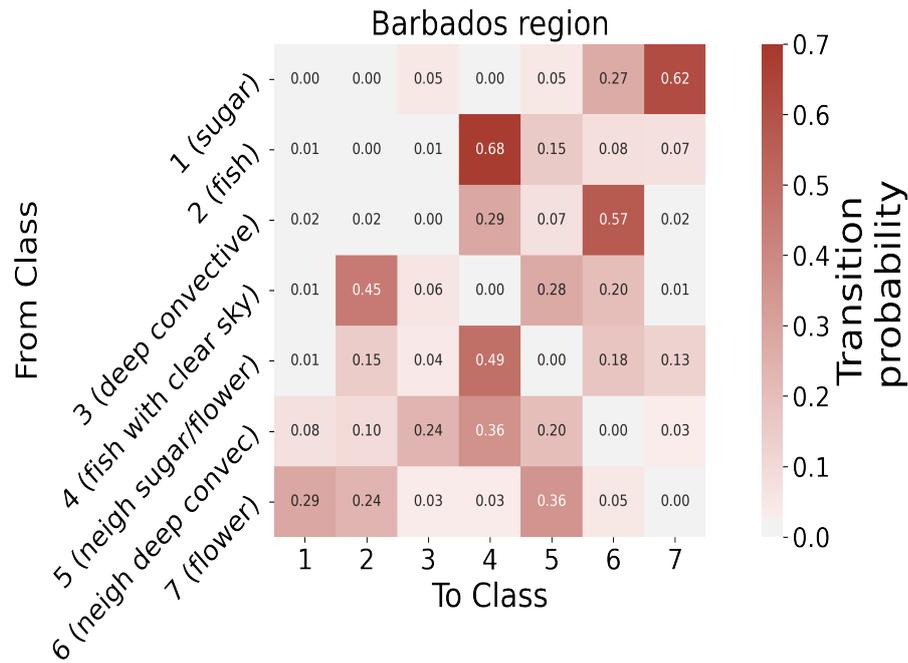


Figure S7. Matrix representing the probability of transition of each machine identified cloud regime to another over the Barbados cloud observatory. The total number of samples for this analysis are 5,470.

The goal is to understand how close classes should be in the feature space to make a transition and derive a transition probability from one cloud regime to another. We observe that sugar has the highest possibility (62%) of transforming to flower-type cloud systems (Figure reviewer-2:3). A mature developed fish could become a fish with more open sky with 68% probability (possibly

due to cold pool downdraft effect). The deep convective systems follow their neighboring class in the continuum, i.e., class six. Fish with more open sky is followed by mature fish-type structures. Class five, which has cloud cells typically found in sugar, gravel, and flower, has 49% chance of a follow-up by fish with more open sky. Finally, the flower class has 36% chance of a transition to class 5 or 30% chance to sugar type distributions.

References

- Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001). On the surprising behavior of distance metrics in high dimensional space. In J. Van den Bussche & V. Vianu (Eds.), *Database theory — icdt 2001* (pp. 420–434). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Bottou, L. (2012). Stochastic gradient descent tricks. In G. Montavon, G. B. Orr, & K.-R. Müller (Eds.), *Neural networks: Tricks of the trade: Second edition* (pp. 421–436). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from https://doi.org/10.1007/978-3-642-35289-8_25 doi: 10.1007/978-3-642-35289-8_25
- Bridle, J. S. (1989). Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In *Nato neurocomputing*. Retrieved from <https://api.semanticscholar.org/CorpusID:59636530>
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). *Emerging properties in self-supervised vision transformers*.
- Chatterjee, D., Acquistapace, C., Deneke, H., & Crewell, S. (2023). Understanding cloud systems structure and organization using a machine’s self-learning approach. *Artificial Intelligence for the Earth Systems*. Retrieved from <https://journals.ametsoc.org/view/journals/aies/aop/AIES-D-22-0096.1/AIES-D-22-0096.1.xml> doi: <https://doi.org/>

10.1175/AIES-D-22-0096.1

Denby, L. (2020). Discovering the Importance of Mesoscale Cloud Organization Through Unsupervised Classification. *Geophysical Research Letters*, *47*(1), e2019GL085190. doi: 10.1029/2019GL085190

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... Houlsby, N. (2021). *An image is worth 16x16 words: Transformers for image recognition at scale.*

Fukushima, K. (1975). Cognitron: A self-organizing multilayered neural network. *Biological Cybernetics*, *20*(3-4), 121–136. Retrieved 2022-08-30, from <http://link.springer.com/10.1007/BF00342633> doi: 10.1007/BF00342633

Gao, J., Li, F., Wang, B., & Liang, H. (2021). Unsupervised nonlinear adaptive manifold learning for global and local information. *Tsinghua Science and Technology*, *26*(2), 163-171. doi: 10.26599/TST.2019.9010049

Goyal, P., Duval, Q., Reizenstein, J., Leavitt, M., Xu, M., Lefaudeaux, B., ... Misra, I. (2021). *Vissl*. <https://github.com/facebookresearch/vissl>.

He, K., Zhang, X., Ren, S., & Sun, J. (2015, December). Deep Residual Learning for Image Recognition. Retrieved 2022-08-30, from <http://arxiv.org/abs/1512.03385> (arXiv:1512.03385 [cs])

Hendrycks, D., & Gimpel, K. (2023). *Gaussian error linear units (gelus)*. Retrieved from <https://arxiv.org/abs/1606.08415>

Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., & Shah, M. (2022, jan). Transformers in vision: A survey. *ACM Computing Surveys*, *54*(10s), 1–41. doi: 10.1145/3505244

- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986, October). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536. Retrieved from <https://doi.org/10.1038/323533a0> doi: 10.1038/323533a0
- Tenenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, *290*(5500), 2319–2323. Retrieved from <https://www.science.org/doi/abs/10.1126/science.290.5500.2319> doi: 10.1126/science.290.5500.2319
- van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, *9*(86), 2579–2605. Retrieved from <http://jmlr.org/papers/v9/vandermaaten08a.html>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2023). *Attention is all you need*. Retrieved from <https://arxiv.org/abs/1706.03762>