https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

Investigating the cloud liquid water and droplet concentration relationship across cloud morphologies with machine learning

Andrew Geiss, Johannes Mülmenstädt, Adam C. Varble, Matthew W. Christensen, and Heng Xiao
Pacific Northwest National Laboratory, Richland WA, USA
Correspondence: Andrew Geiss (andrew.geiss@pnnl.gov)

Abstract. The relationship between cloud droplet number concentration (N_d) and cloud liquid water path ($\mathcal{L}$) is investigated in the context of detailed information about cloud morphology derived from self-supervised machine learning (ML) interpretation of satellite imagery over the North Atlantic ocean. While different cloud morphologies occupy distinct regions in the $\mathcal{L}$-N_d phase space, the "inverted-V" relationship observed between $\mathcal{L}$ and N_d does not arise as a result of their distribution in this space, and most cloud morphologies independently exhibit the inverted-V pattern to some degree. A novel approach to investigate this relationship is demonstrated using continuous vector space image embedding representations of cloud morphology produced by the self-supervised ML. Analysis of non-linear regression between these cloud morphology representations, N_d, and $\mathcal{L}$ indicates that while cloud morphology information can explain most of the variance in $\mathcal{L}$ between cloud scenes, the portion of $\mathcal{L}$ variance that can be explained by N_d is mostly independent of cloud morphology and explains the inverted-V pattern.

1 Introduction

Clouds significantly alter the Earth's radiative budget, influencing both weather and climate. Cloud radiative effects are determined by a combination of cloud microphysical (e.g. hydrometeor size and phase), macrophysical (e.g. cloud albedo and liquid water path), and morphological properties (e.g. organization and cloud fraction). Marine clouds have a particularly large impact on shortwave radiation – they are bright objects over a dark surface and reflect sunlight back to space that would otherwise be absorbed at the surface. Aerosols act as cloud condensation nuclei (CCN), and influence the microphysical properties of clouds. Increasing the number of CCN, and thus the droplet number concentration in a cloud, without changing the amount of water present, increases cloud albedo (Twomey, 1977). Furthermore, reduced droplet size and increased number concentration in a warm marine cloud can suppress precipitation, leading to longer cloud lifetimes (Albrecht, 1989). This aerosol cloud interaction (ACI) (Fan et al., 2016), remains one of the largest sources of internal uncertainty in climate projections (Boucher et al., 2013; Bellouin et al., 2020; Forster et al., 2021). A comprehensive understanding of ACI would enable improved representation in climate and weather simulations, but has remained elusive for a myriad of reasons: clouds respond in multiple ways to

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

changes in aerosol loading and isolating these responses can be difficult; satellite remote sensing observations of clouds are plentiful but have limitations depending on the instrument and retrieval; in-situ observations of marine clouds (with aircraft for example) are sparse and expensive; truly co-located remote sensing observations of both aerosol and cloud are not possible; the conditions of the lower atmosphere also have a substantial impact on marine cloud properties that can co-vary with ACI; and finally, there are many different types of marine cloud (Wood, 2012) that might respond in different ways to changes in aerosol loading. This work focuses on the last of these issues, and attempts to untangle the co-variance of observed cloud-aerosol interactions with different marine cloud morphologies.

The relationship between cloud liquid water path ($\mathcal{L}$) and droplet number concentration (N_d) is key to ACI. Changes in N_d will change cloud albedo, altering clouds' radiative properties, but can also cause increases or decreases in $\mathcal{L}$ that enhance or suppress this effect. Constraining this relationship is important to understanding the response of the Earth's climate to changes in atmospheric aerosols, but the results of past studies of the $\mathcal{L}$-N_d relationship have varied widely depending on region, cloud type, data sources, and methodology (Han et al., 2002; Wang et al., 2003; Bretherton et al., 2007; Quaas et al., 2009; Chen et al., 2014; Gryspeerdt et al., 2014; Grosvenor et al., 2017; Bender et al., 2019; Glassmeier et al., 2021; Arola et al., 2022). Gryspeerdt et al. (2019) identified that when satellite retrieved cloud properties are aggregated, an "inverted-V" pattern emerges (some examples can be seen in Figure 4 of this manuscript). Specifically, the inverted-V appears in the conditional probability density function $P(\mathcal{L}|N_d)$ plotted in log coordinates, and shows that $\mathcal{L}$ tends to rapidly increase with increasing N_d for low-N_d clouds ($N_d <\sim 20-30cm^{-3}$), and slowly decrease with increasing N_d for high-N_d clouds (Gryspeerdt et al., 2019; Mülmenstädt et al., 2024; Goren et al., 2025). This pattern has been hypothesized to result from a combination of physical effects: in the ascending leg of the V, cloud droplets are sparse and large, and clouds with these properties often precipitate. Increasing N_d leads to smaller and more numerous droplets, which reduces collision-coalescence and suppresses precipitation, increasing $\mathcal{L}$ (Albrecht, 1989). In the descending branch, increasing N_d and reducing droplet size enhances the rate at which droplets evaporate (Pruppacher and Klett, 1997; Xue and Feingold, 2006) and smaller droplets experience less sedimentation, which enhances cloud top radiative cooling and entrainment, further increasing droplet evaporation and reducing $\mathcal{L}$ (Ackerman et al., 2004; Bretherton et al., 2007). The existence of this pattern complicates efforts to understand the influence of altered aerosol loading on climate, because it implies that the radiative properties of different clouds will respond to changes in aerosol loading in different ways.

There has recently been significant effort to understand the inverted-V pattern. Gryspeerdt et al. (2019) analyzed the $\mathcal{L}$-N_d relationship over the world's oceans using multiple instruments (AMSR-E and MODIS) in an effort to constrain N_d controls on $\mathcal{L}$ and to reconcile past studies indicating different slopes between these two variables ($\frac{\mathrm{dln}(\mathcal{L})}{\mathrm{dln}(N_d)}$), and showed that this inverted-V pattern is robustly observed across different instruments and regions though the slopes of the two legs of the inverted-V vary. Mülmenstädt et al. (2024) compared observed $\mathcal{L}$-N_d slopes to those produced by climate models, and showed that while CMIP5-era models generally did not reproduce the negative slope in high N_d conditions, CMIP6-era models can, though still do not agree well with observations on the exact shape and magnitude of the $\mathcal{L}$-N_d relationship. Most closely related to this study is Goren et al. (2025), who showed evidence that the inverted-V largely arises from co-variability of $\mathcal{L}$ and N_d in response to different environmental and microphysical controls. They also showed that marine stratocumulus types play an

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

important role, with open cellular convection occupying the low-N_d portion of the V and closed cellular convection occupying the high N_d portion of the V, indicating that different marine cloud morphologies likely have different responses to changes in background aerosol concentration. Goren et al. (2026) used a ternary encoding of cloud optical depth as a continuous proxy for cloud morphology by binning the fraction of pixels in 2° cloud scenes with cloud optical depth in three different ranges. They found that almost all morphologies defined this way show negative $\mathcal{L}$ susceptibility to N_d and that the sign and magnitude of cloud albedo susceptibility to N_d is highly dependent on cloud type. The work by Goren et al. (2025) and Goren et al. (2026) suggests the inverted-V arises, at least partially, as an artifact of aggregating multiple cloud types with different $\mathcal{L}$-N_d properties.

Marine stratocumulus cloud organizes in a variety of distinct morphological patterns (Wood, 2012; Houze, 2014). These patterns are easily recognizable in visible satellite imagery, and have been heavily studied because of their different radiative and microphysical properties that affect weather and climate. There is now a vast amount of cloud imagery from Earth observing satellites, but these datasets are far too large to be manually labeled by researchers, so machine learning (ML) has emerged as key tool in studying cloud morphology. Wood and Hartmann (2006) developed a 4-class fully-connected feed-forward neural network-based (Bishop, 2006) satellite imagery classifier to study closed-cell and open-cell convection, two important marine stratocumulus regimes. Rasp et al. (2020) created the largest hand labeled marine cloud imagery dataset to date using crowd-sourcing, and trained a deep convolutional neural network (CNN) (Goodfellow et al., 2016) segmentation algorithm to differentiate between the four subtropical cloud types studied by Stevens et al. (2020). Segal Rozenhaimer et al. (2023) designed a CNN segmentation algorithm using the same classes as Wood and Hartmann (2006). Finally, Yuan et al. (2020) developed a supervised CNN using a smaller training dataset but including six cloud categories. A drawback of supervised image recognition schemes is that they rely on large hand-labeled training datasets (e.g. Deng et al. (2009)). While hand-generating these training sets is possible (whereas labeling a full satellite climatology is not), it is still extremely time consuming and typically requires labelers with expertise in the field. Supervised learning schemes also require all categories of interest to be known before training. More recently, self-supervised machine learning algorithms have been applied to this problem (Yuan, 2019; Denby, 2020; Kurihana et al., 2019, 2022; Geiss et al., 2024). These algorithms can train on large unlabeled imagery datasets, and do not rely on extensive human labeling. They are generally built to map images to lower dimensional vector embeddings that are representative of the semantic content of the image, and while they do not directly perform classification, they can be applied to classification tasks without the extensive person hours required to generate training data for supervised algorithms (Chen et al., 2016; Zbontar et al., 2021; Chen et al., 2020).

This work further investigates the relationship between marine cloud morphology and the inverted-V $\mathcal{L}$-N_d pattern. Here, to assess cloud morphology, we leverage the approach developed by Geiss et al. (2024) based on Zbontar et al. (2021). The algorithm is a CNN that maps satellite marine cloud images to lower dimensional vector embeddings that are representative of cloud morphology. We apply the algorithm to 10-years of satellite cloud images from the North Atlantic and assess the $\mathcal{L}$-N_d relationship across different cloud morphologies both using an unsupervised clustering approach and by interrogating neural networks trained to represent the non-linear mappings between the cloud imagery vector embeddings and $\mathcal{L}$ observations. Our primary finding is that while different marine cloud types occupy different regions of $\mathcal{L}$-N_d space, the inverted-V pattern

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

appears consistently across multiple cloud types, and remains when we control for the relationship between cloud morphology and $\mathcal{L}$ using machine learning. Even so, the slopes of the branches of the inverted-V vary between different cloud morphologies, and this should be considered when assessing ACI.

2 Data

This study uses cloud imagery and microphysical retrievals collected by the MODerate resolution Imaging Spectroradiometer (MODIS) (Platnick et al., 2017). MODIS is deployed on two polar orbiting satellites "Terra" and "Aqua," and here, we use co-located Terra MODIS collection 6.1 level 1 (L1) reflectance imagery and level 2 (L2) cloud property retrievals (more details are provided below), which are packaged as separate datasets, though they are collected simultaneously. The L1 imagery is used to assess cloud morphology, while the L2 data are used for co-located retrievals of $\mathcal{L}$ and N_d. The analysis is focused on the North Atlantic region, with the intent of providing results that can later be compared to detailed in-situ and ground-based remote sensing observations from the Atmospheric Radiation Measurement East North Atlantic (ARM ENA) facility, located in the Azores (Mather and Voyles, 2013). MODIS data are archived in "granule" files, which contain 1350 km (across-swath) by 2030 km (along-swath) images of the Earth at 500-m resolution for the L1 images and 1-km resolution (at nadir) for the various L2 cloud retrievals used here. The dataset used here consists of all MODIS granules that at least partially pass within $\pm 20°$ latitude and $\pm 20°$ longitude of the ENA site (located at approximately 39°N and 28°W) between July 1st 2013 and January 1st 2023, which is the majority of the period that site has been deployed and provides good coverage of the ENA region. Figure 1a shows the spatial distribution of all of the MODIS data retrieved for this study, and a list of all the MODIS granules used is provided in the data supplement.

2.1 MODIS Level-1 reflectance imagery

This work uses the self-supervised algorithm developed by Geiss et al. (2024) to generate vector embeddings of MODIS L1 reflectance imagery that are representative of the cloud morphology contained in the images. The embeddings are then used to perform automated clustering of cloud imagery into different morphological regimes, and to perform non-linear regression to represent the relationship between vector embeddings of cloud morphology and $\mathcal{L}$. The L1 imagery is from the Terra MODIS version 6.1 level-1B half-kilometer product (MODIS Characterization Support Team (2017), file names: "MOD02HKM", see the data availability section). The ML algorithm uses single channel 550-nm reflectance images (variable name "EV_500_RefSB"). It operates on 256x256-pixel samples (referred to here as "image-chips") that are taken from the granule file with a stride of 128-pixels. The granule files are center-cropped such that the resulting width and height are multiples of 128 pixels. 10.8×10^6 image chips were retrieved in total. The focus of this study is low marine clouds, and we apply some simple filters to ensure that the images included in the unsupervised clustering process, described below, are well illuminated and are not dominated by land, high cloud, or clear-sky pixels. These constraints are applied using averages of the MODIS L2 product aggregated over each image chip:

– Less than 10% land cover.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

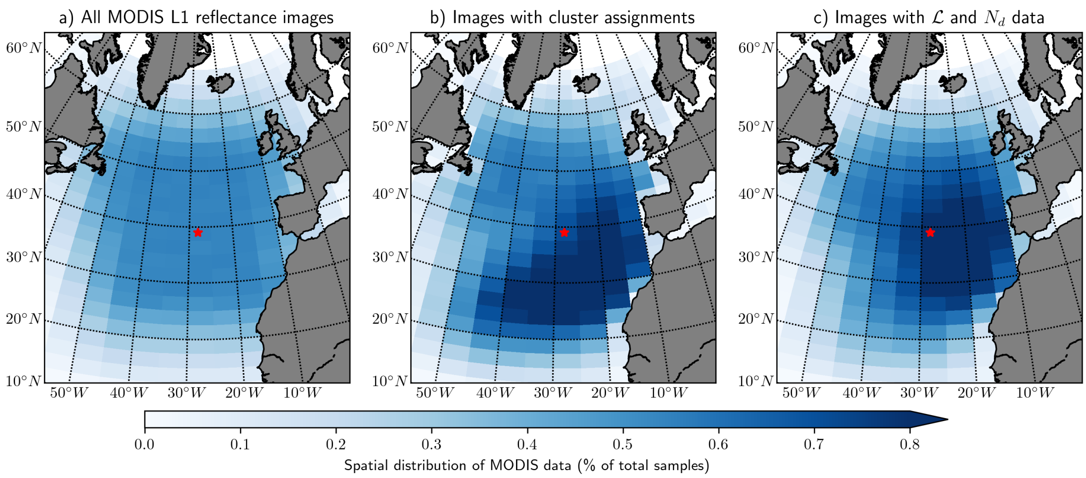

Figure 1. Panel a) the spatial distribution of all MODIS imagery retrieved for use in this study (10.8×10^6 image chips). Panel b) the spatial distribution of all samples that met the criteria for generating cloud morphology labels described in Section 2.1 (3.7×10^6 image chips). Panel c) the spatial distribution of all samples that have cloud morphology labels and valid $\mathcal{L}$ and N_d retrievals based on the criteria in Section 2.2 (3.4×10^6 image chips). The samples in panel c are a subset of those in panel b which are a subset of those in panel a. The star in the center of each panel is the location of the ARM ENA site. The spatial distributions are expressed as the percentage of total samples falling into $5^\circ \times 5^\circ$ grid-cells.

- Solar zenith angle less than 70°.
- Greater than 20% cloud fraction.
- Less than 5% high cloud (cloud top pressure < 500-hPa) pixels.

3.7×10^6 total image chips meet these criteria and the resulting spatial density of samples is shown in Figure 1b. These simple constraints are applied for the k-mean clustering process, described below, to ensure that the clusters are not dominated by scenes irrelevant to the study, but much stricter filtering, described in the next section, is applied for analysis involving cloud $\mathcal{L}$ and N_d.

2.2 MODIS Level-2 cloud product

We also use co-located data from the Terra MODIS collection 6.1 L2 cloud product (Platnick et al. (2015), file names: "MOD06_L2", see the data availability section) which includes numerous passive cloud property retrievals at resolutions ranging from 1-5 km. $\mathcal{L}$, cloud optical depth (τ_c), and cloud droplet effective radius (r_e) retrievals are packaged with this dataset at 1-km resolution, but N_d estimates are not. We calculate N_d following Grosvenor et al. (2018) and Mülmenstädt et al. (2024) with:

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

$$N_d = \frac{\sqrt{5}}{2\pi k}\left(\frac{f_{\mathrm{ad}} c_{\mathrm{w}} \tau_{\mathrm{c}}}{Q_{\mathrm{ext}} \rho_{\mathrm{w}} r_{\mathrm{e}}^5}\right)^{\frac{1}{2}} \quad (1)$$

where the terms $f_{\mathrm{ad}} = 1$ (dimensionless) is the sub-adiabatic factor, $c_{\mathrm{w}} = 2\times10^{-6}$ (kg m^{-4}) is the condensation rate, $Q_{\mathrm{ext}} = 2$ (dimensionless) is the extinction efficiency, $\rho_{\mathrm{w}} = 10^3$ (kg m^{-3}) is the density of liquid water, and $k = 1$ (dimensionless) is the cubed ratio between the volumetric mean droplet radius and the effective radius. These are each assumed to be constant valued.

Past studies using passive satellite $\mathcal{L}$ and N_d retrievals have applied strict data quality constraints, and we apply filtering based on Grosvenor et al. (2018), Gryspeerdt et al. (2019), and Goren et al. (2025). Specifically, we apply the following constraints at a 1-km pixel level:

- Single-layer cloud, according to the MODIS multi-layer cloud flag (“Cloud_Multi_Layer_Flag” = 1).
- Liquid phase, according to the cloud phase flag (“Cloud_Phase_Optical_Properties” = 2)
- Cloud top pressure greater than 600 hPa.
- Solar zenith angle less than 65°.
- Viewing zenith angle less than 55 °.
- Cloud top temperature greater than 270 K.
- Cloud optical depth greater than 1.
- Cloud droplet effective radius greater than 3 μm
- Sub-pixel cloud variance index less than 30 (“Cloud_Mask_SPI” < 30).

These constraints address some of the main sources of uncertainty in the cloud property retrievals and ensure that their underlying assumptions are correct where possible. MODIS-derived $\mathcal{L}$ and N_d rely on accurate passive retrievals of τ_c and r_e, which are themselves limited (Merk et al., 2016; Grosvenor et al., 2018). The retrievals assume spherical liquid droplets and the presence of ice crystals can lead to overestimates of $\mathcal{L}$, hence the cloud top temperature and MODIS phase flag constraints. Bi-spectral retrievals of $\mathcal{L}$ and N_d are limited due to cloud heterogeneity (Arola et al., 2022), failure of retrieval algorithms to account for 3D radiative effects (Loveridge and Di Girolamo, 2024), and have greater uncertainty for thin and broken cloud (Zhang and Platnick, 2011), motivating the restrictions on sub-pixel variance and τ_c. N_d biases have been observed for high solar zenith angles and sensor viewing angle (Grosvenor et al., 2018), so these are also constrained. There are additional sources of uncertainty that are not as easily accounted for, and results based on MODIS L2 cloud retrievals should be interpreted with this in mind. For example, the assumption of constant adiabaticity in Eq. 1 or the fact that cloud $\mathcal{L}$ and N_d retrievals both rely on τ_c and r_e retrievals and thus errors in those retrievals can lead to correlated errors in $\mathcal{L}$ and N_d (Gryspeerdt et al., 2019).

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

Despite these concerns, past work has shown that with appropriate filtering these retrievals are reliable (Painemal et al., 2025), and they have previously been used to study the $\mathcal{L}$-N_d relationship (Gryspeerdt et al., 2019; Goren et al., 2025).

In the following sections, we perform analysis of these data within the image chips (256x256-pixel L1 and co-located 128x128-pixel L2 sub-samples of the MODIS granules). To compute the average $\mathcal{L}$ and N_d values corresponding to an L1 image chip, we average only valid pixels according to the above constraints and report the averages only for chips with at least 10% valid pixels (resulting in the average value for cloudy pixels not the area average over the full image). Figure 1c shows the geographical distribution of image chips that have usable $\mathcal{L}$ and N_d data. The final dataset consists of approximately 3.4×10^6 valid image chips representing 10-years of North Atlantic warm marine cloud observations.

3 Machine learning of cloud morphology

Geiss et al. (2024) developed a self-supervised CNN that maps satellite cloud images to vector embeddings that represent the cloud morphology present in the image. The CNN was trained on one year (2020) of global MODIS Terra 500-m resolution band 4 (550-nm) reflectance images of cloudy scenes over the world's oceans. It maps 256x256-pixel images to 1024-value vector embeddings and it was trained using the "Barlow Twins" algorithm proposed by (Zbontar et al., 2021). The method relies on finding pairs of images with similar semantic content to avoid requiring human labels and leverages the fact that cloud morphology is spatially autocorrelated to generate training samples. We have used the pre-trained CNN available from Geiss (2023) to map the MODIS image chips discussed in Section 2.1 to vector embeddings. Geiss et al. (2024) showed that, with an appropriate distance metric, images can be compared in this embedding space to reliably determine the similarity between the cloud morphologies shown in the images.

3.1 Cloud type clustering

The k-means algorithm (Hartigan, 1975; Bishop, 2006) was then used to perform unsupervised clustering of the ENA imagery vector embeddings to develop a catalog of common cloud morphologies in the vicinity of the ENA site. K-means was selected after a brief comparison to a hierarchal clustering algorithm similar to the algorithm used by Kurihana et al. (2022). Within our image vector embedding space, we found that the k-means algorithm provided better separation between cloud categories when applied to existing labeled cloud imagery datasets (Yuan et al., 2020; Rasp et al., 2020) (these results are not shown). Following Geiss et al. (2024), the image vector embeddings are best compared using the cosine similarity metric:

$$\text{Cosine Similarity} = \frac{\mathbf{x} \cdot \mathbf{c}}{||\mathbf{x}|| \, ||\mathbf{c}||} \tag{2}$$

which represents the similarity of a sample $(\mathbf{x})$ to its assigned cluster center $(\mathbf{c})$. To maximize the cosine similarity between each sample and its assigned cluster center, we customized the k-means algorithm to optimize:

$$\frac{1}{N}\sum_{n \in N} \max_k \left\{ \frac{\mathbf{x}_n \cdot \mathbf{c}_k}{||\mathbf{x}_n|| \, ||\mathbf{c}_k||} : k \in K \right\} \tag{3}$$

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

where n iterates over the N total samples and k iterates over the K cluster centers.

The k-means algorithm does not scale well to large datasets and has complexity proportional to $\mathcal{O}(NDKI)$, where $N \times D$ are the dimensions of the vector embedding dataset, K is the number of clusters used, and I is the number of iterations required for convergence of the algorithm, which is dependent on N and K. However, the algorithm can be accelerated with an optimal initial choice of cluster centers (Celebi et al., 2013), and we employed a method similar to Bradley et al. (1998). First we applied the k-means algorithm separately to ten randomly selected 1% subsets of the vector embeddings. Then the set of cluster centers from the subset that yielded the best final k-means distance metric (Eq. 3) were used to initialize the k-means algorithm when it was run on the full dataset. In all cases, the algorithm was run until it fully converged.

A notable drawback of k-means is that it requires choosing the number of clusters (K) before the clustering is performed. This means that 1) it is limited to finding exactly K clusters even if there are a different number of natural clusters present in the data distribution, splitting natural clusters when K is too high, and 2) it has a tendency to ignore minority clusters when they are small enough. So while the algorithm is generally considered to be unsupervised, it requires some expert input regarding the choice of K. We found through manual examination of the resulting clusters, that for large values of K (128), the most common cloud regimes were often split between many clusters. Meanwhile, for small values of K (16), less common, but clearly distinct cloud types were ignored, and grouped into clusters dominated by a more common cloud type. As a compromise, we chose to run k-means with a large number of clusters ($K = 64$), and manually group together redundant clusters. Ultimately, this process resulted in 12 hand-selected super-clusters, each composed of multiple k-means clusters. This choice removes some of the objectivity from the self-supervised image processing and unsupervised clustering, but most of the process is still done without human input, and the approach significantly cuts down on the number of person-hours that would be required to achieve a similar result using supervised algorithms and hand-labeling of individual cloud images.

3.2 Overview of North Atlantic cloud regimes

The clustering procedure generated 12 total low cloud morphological regimes (super-clusters) from the ENA region. Each of these is comprised of multiple clusters from the k-means algorithm. Of the 12 super-clusters, around 8 seem to represent truly unique cloud-types, 2 appear to represent cases with multiple cloud types present, and 2 are redundant, with visually different images but similar cloud types to another category. The k-means algorithm was also able to isolate some cases that should not have been included in the analysis but were not excluded by the filters applied to the imagery data (Section 2.1). For instance, not all cases with sea ice were removed, but because its spatial structure easily differentiates it from cloud, the k-means algorithm grouped these cases in a single cluster and they were discarded. k-means also found some seemingly clear-sky cases, with no visible cloud (likely there is some very optically thin cloud present in these cases), and this cluster was also discarded. Figure 2 shows examples of 4 images from one of the k-means clusters selected from each of the 12 super-clusters. The samples are selected randomly from those within the 10% of samples that are closest to their assigned cluster center, because the samples close to the edges of the clusters are often a much poorer match for the cloud morphology represented by that cluster. Overall, we primarily identified clusters that are a good match for the unique low marine cloud regimes identified in past meteorological studies. Below, we list each of the super-clusters with short descriptions, and include in parenthesis the

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

number of k-means clusters included in each super cluster and the total percentage of the cloud imagery dataset that has been assigned to each regime. The supplementary material also includes a collection of figures showing examples of each of the 64 k-means clusters and information about their spatial and temporal distribution and $\mathcal{L}$-N_d properties.

a) **Open cellular convection:** (8 clusters, 11.9%) This is a common cloud regime over the extra-tropical oceans (Wood, 2012; Goren et al., 2025). It is characterized by a honeycomb cellular structure with clear areas of descending air in the cell centers surrounded by inter-connected cell edge regions of convective cloud. These clouds frequently occur near the ENA site.

b) **Closed cellular convection:** (6 clusters, 8.2%) Closed cellular convection is a stratocumulus regime composed of distinct cells with strong rising motion near the center surrounded by a much larger region of weaker subsidence (Wood, 2012; Goren et al., 2025). Regions of closed cellular convection typically have nearly 100% cloud cover.

c) **Flowers:** (3 clusters, 5.2%) This is a primarily subtropical cloud regime called “flowers” by Stevens et al. (2020) and also called “disorganized convection” by (Yuan et al., 2020). These are large (20-200 km wide) clouds characterized by a convective core that is surrounded by thinner cloud close to the top of the marine boundary layer. Individual “flowers” are typically separated from each other by clear areas.

d) **Sugar:** (6 clusters, 9.5%) These samples are composed of many small, isolated, and disorganized (meaning they do not form a clear pattern or repeating large-scale structure) cloud elements. They are primarily composed of cumulus humilis, but many samples contain some isolated larger, and potentially deeper cloud elements. "Sugar" is borrowed from Stevens et al. (2020), and (Yuan et al., 2020) refers to these cases as “suppressed cumulus.”

e) **Subtropical stratocumulus:** (4 clusters, 5.0%) This category contains images of stratocumulus decks that completely or almost completely cover a scene. The self-supervised algorithm can visually differentiate between stratocumulus decks that occur in the subtropics and those in the mid-latitude storm track regions and these are separated into two categories. This type of marine cloud commonly occurs on the eastern boundaries of the world’s ocean basins where the descending branch of the Hadley cell, cool ocean surface temperatures, and relatively little synoptic variability combine to form marine boundary layers that are optimal for a persistent stratus layer (Wood, 2012). Small scale repeating patterns of convective elements are often evident in the images, but are not consistent, large, or thick enough to be closed cellular convection.

f) **Storm-track stratocumulus** (6 clusters, 8.1%) These scenes are visually similar to the subtropical stratocumulus category but tend to occur in the mid-latitude storm tracks. They typically appear to be brighter, thicker cloud, but are also stratocumulus decks with a clear repeating texture from the smaller cumulus elements they are composed of, which are too small to be considered closed cellular convection.

g) **Thick stratus** (5 clusters, 10.0%) These scenes contain low-level stratiform cloud that does not have a clear repeating pattern or texture. These scenes have nearly complete cloud coverage.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

h) **Broken stratus:** (6 clusters, 12.2%) These are images showing the transition between stratus decks and clear sky regions. The k-means algorithm likely has separated these into a cluster because the images are visually very different from the other stratiform categories even though the cloud types are similar.

i) **Semi-organized convection:** (3 clusters, 5.4%) This super-cluster contains cases primarily composed of convective elements that have a vague spatial repetition, but are not organized enough to be confidently called cellular convection. The convective elements are visually similar to the open cellular convection category, meaning they are bright, presumably thick, cloud surrounded by clear-sky regions.

j) **Clustered cumulus:** (6 clusters, 8.2%) These cases contain sparse, bright, moderately sized (5-20 km), cumulus elements that are not organized into any obvious large-scale pattern. They occur in both the mid-latitude storm track regions and the subtropics. Yuan et al. (2020) calls these "clustered cumulus."

k) **Mixed clustered cumulus and stratiform** (4 clusters, 7.3%) These scenes contain both stratiform cloud regions and disorganized convective elements. They do not have clear repeating spatial patterns, but are common enough to warrant their own category.

l) **Mixed sugar and clustered cumulus** (5 clusters, 9.1%) These cases contain both cumulus humilis and much larger disorganized convective elements in the same scene (the rough linear organization of the smaller clouds in the examples in Figure 2 are a feature of the sub-cluster shown, not the super-cluster).

Sections 4.1 and 4.2 use these clusters to investigate the $\mathcal{L}$-N_d relationship in the context of discrete cloud morphology regimes.

3.3 A continuous representation of cloud type

The super-clusters described above largely align with the various marine cloud types that have been identified and studied by past authors (Stevens et al., 2020; Yuan et al., 2020; Wood, 2012; Houze, 2014). In reality, there are a range of different cloud organizational patterns, and while many canonical examples of the frequently studied organizational patterns in the literature (e.g. closed/open cellular convection) can be unambiguously identified in satellite imagery, there are many other cases that are not as well defined. This is reinforced by examination of the the 64 clusters from k-means (S1). Even within each of the super-clusters identified above, there is substantial variability in the appearance of the cloud type across each of the k-means clusters composing it. There are also many cloud scenes that contain multiple cloud types, which in some cases were common enough to warrant their own super-cluster in Section 3.2, or might exist somewhere on the boundary between two different cloud morphologies. This will be unsurprising to anyone who has spent a significant amount of time working with satellite cloud imagery – often-times transitions between cloud types can be seen occurring smoothly over a region or time period, and it is not clear where to draw a distinction. For example, the open cellular convection and the semi-organized convection categories discussed above are closely related and the boundary between the two is often ambiguous.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

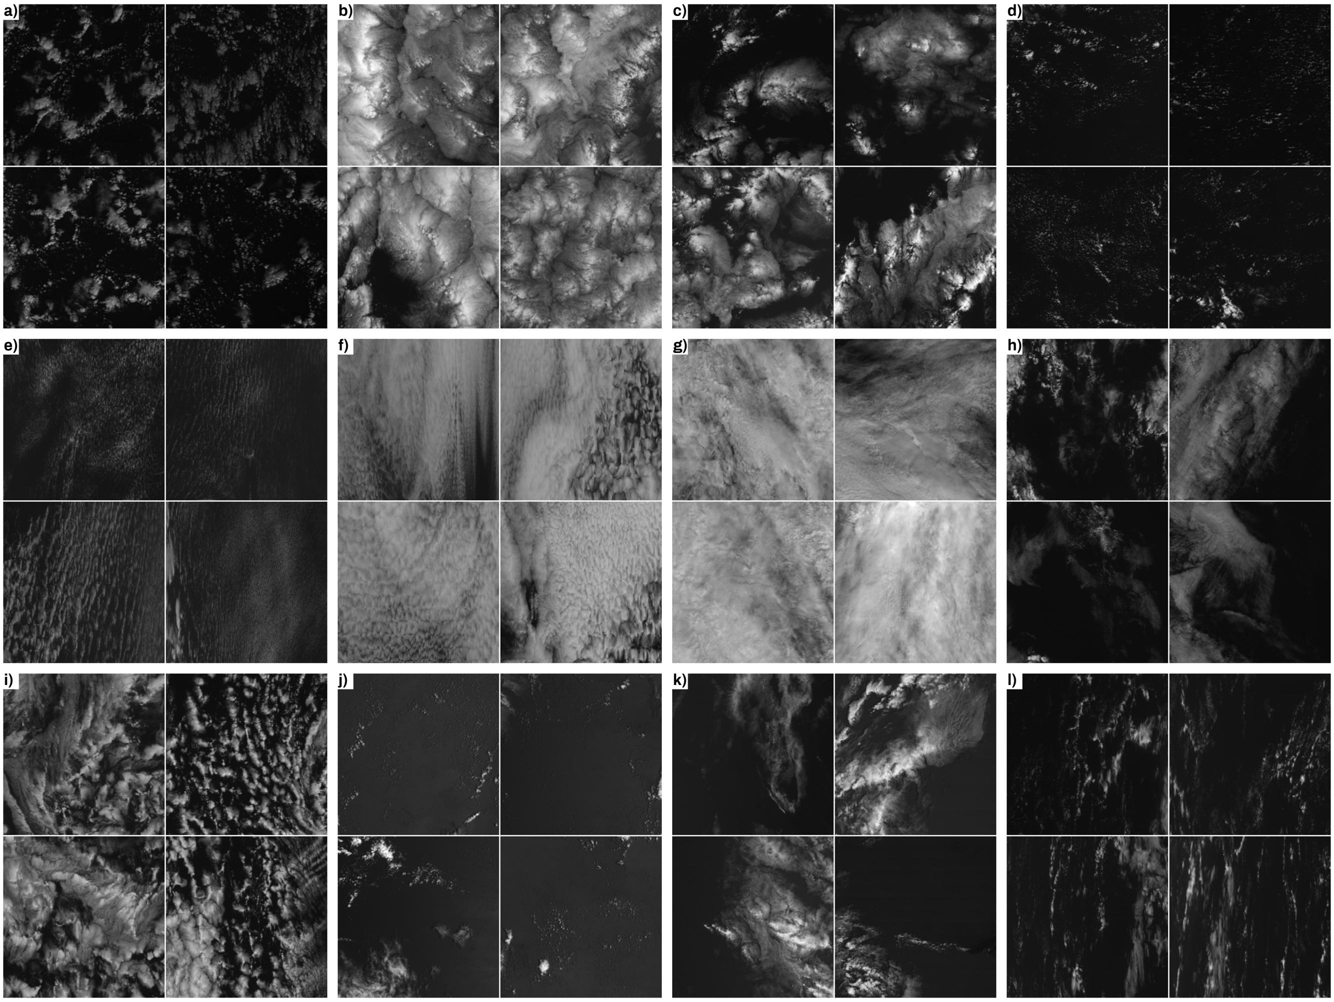

Figure 2. Examples of each of the cloud morphology clusters described in Section 3.2. Each panel contains four randomly selected sample images from the 10% of cases closest to the corresponding cluster center (to ensure an unambiguous representation of that cluster's cloud type). The cluster names are a) open cellular convection, b) closed cellular convection, c) flowers, d) sugar, e) subtropical stratocumulus, f) storm-track stratocumulus, g) thick stratus, h) broken stratus, i) semi-organized convection, j) clustered cumulus, k) mixed clustered cumulus and stratiform, l) mixed sugar and clustered cumulus.

Dividing cloud into a limited number of well-defined morphologies helps simplify an extremely complex natural system to something that is more tractable for cloud studies. While these classifications are a useful tool, ML approaches do not
295 necessarily need this type of simplification. The vector embeddings of cloud imagery from the self-supervised algorithm map cloud imagery to a high dimensional *continuous* space, that provides a unique tool for investigating cloud morphology. In Sections 4.3 and 4.4 we present two experiments that involve using simple, multi-layer perceptron (Bishop, 2006) type neural networks (NNs) as generic non-linear function fitting tools to map between continuous cloud morphology information and

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

cloud physical properties. The NNs can then be interrogated to learn about how cloud type impacts the relationship between cloud $\mathcal{L}$ and N_d. This approach eliminates a major source of human bias from the analysis and no difficult judgments need to be made regarding the total number of cloud types, which cloud types are relevant, or how to label observations.

4 The $\mathcal{L}$-N_d relationship in North Atlantic clouds

This section analyzes the $\mathcal{L}$-N_d relationship in the context of North Atlantic cloud morphology using several different approaches that address the following questions: Section 4.1) does the inverted-V pattern arise from the location of different cloud types in $\mathcal{L}$-N_d space? Section 4.2) what is the $\mathcal{L}$-N_d relationship within different cloud regimes? Section 4.3) how much $\mathcal{L}$ variance can be described by variance in cloud morphology versus variance in N_d? 4.4) what is the $\mathcal{L}$-N_d relationship when we control for variation due to cloud morphology?

4.1 The distribution of cloud types in $\mathcal{L}$-N_d space

Different cloud types have different $\mathcal{L}$ and N_d properties. For example, in areas of open cellular convection, the majority of the cloud is associated with the upward motion occurring at the edges of the cells, while the large areas of descent in the center of the cells are clear. As a result, these scenes will tend to have high $\mathcal{L}$ with relatively low N_d associated with the convective regions, where there are deep, often precipitating, clouds composed of large drops. Meanwhile, scenes with closed cellular convection will tend to have high N_d and more moderate $\mathcal{L}$ because the they are dominated by large areas of thinner cloud composed of smaller drops, with most of the high-$\mathcal{L}$ cloud occurring where there is rising motion through the marine boundary layer at the center of each cell.

Figure 3 shows mean $\mathcal{L}$ and N_d properties of most of the cloud regimes from Section 3.2, with super-clusters indicated by colors and crosses and k-means clusters shown as dots, overlaid on the joint histogram distribution of MODIS $\mathcal{L}$ and N_d observations for the entire dataset (contours). Because points near the edge of the k-means clusters might be poor matches for the corresponding cloud regime, in panel a, all of the data assigned to each cluster are used to determine the marker positions, while in panel b, only the 25% of cases closest (Eq. 2) to their assigned cluster centers are used.

Different cloud types tend to occupy different locations in $\mathcal{L}$-N_d space. These results are in general agreement with Goren et al. (2025), who showed that open cellular convection typically falls in the low-N_d high-$\mathcal{L}$ region of this space and closed cellular convection occupies the high-N_d moderate-$\mathcal{L}$ part of the space. The additional cloud types considered here provide additional granularity. The subtropical stratocumulus category, which here represents very thin stratocumulus decks with fine-scale (on the order of 1-km) cumulus elements, fall in an extremely low-$\mathcal{L}$ and high-N_d part of the plot. Thick stratus occupies the upper right corner, with both high-N_d and high-$\mathcal{L}$. Meanwhile, categories that contain more broken cloud, like clustered cumulus, sugar, and flowers, occupy the central portion of the plot. Many of the relative positions of the cloud types in this plot have intuitive explanations. For example, the cloud elements in the closed cellular convection category and the flowers category are similar, and they appear nearby each-other in $\mathcal{L}$-N_d space. The flowers category tends to have lower $\mathcal{L}$ and slightly lower N_d though. This is likely because the flowers category has large clear-sky regions between cloud elements and these cloud

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

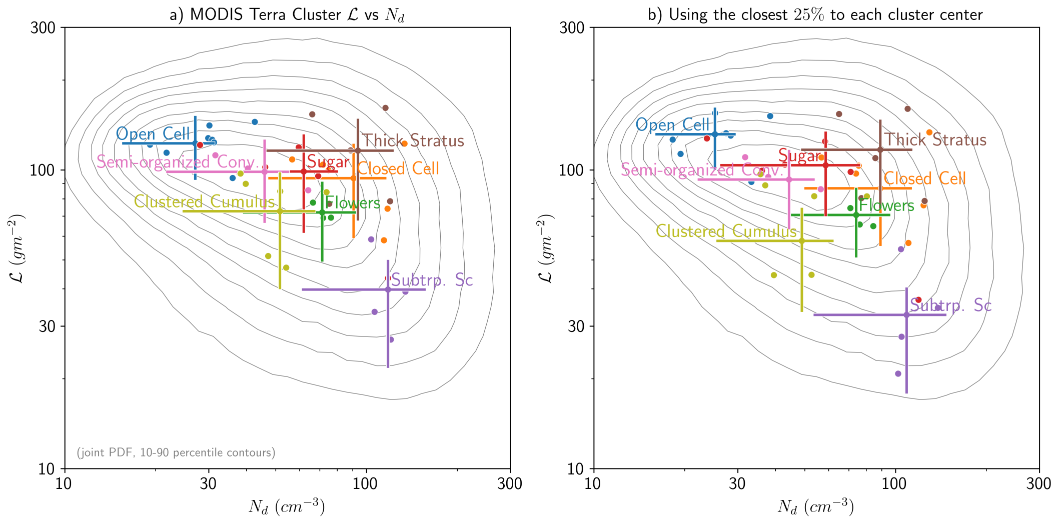

Figure 3. The distribution of cloud-morphology clusters in LWP-Nd space. The coloring represents the human grouped super-clusters, while the dots represent individual k-means clusters. The colored crosses represent the location of the super-cluster means and the inter-quartile range of the samples in the super-cluster. To reduce clutter, we have not included the cloud type groups that represent combinations of different cloud types. Contours show the joint probability distribution of the dataset. Note that the inverted-V pattern is only visible when the conditional probability distribution $P(N_d|\mathcal{L})$ is plotted, and here we are showing the joint distribution of all samples in the space. There are approximately 3.4×10^6 total samples in this dataset, each representing a 128x128 km image, and section 3.2 indicates what fraction of these are assigned to each cluster.

images include more pixels with very thin cloud where there is a transition between cloudy and clear-sky, while in the closed cellular convection scenes, cells are tightly packed together and there is very little thin cloud and clear sky. There is also a clear progression from the open cellular convection category, which occupies the upper left of the space, to semi-organized convection, to clustered cumulus. These cloud types are related, and transitions between these cloud regimes over a large spatial region can frequently be seen in satellite images. It is unsurprising that as samples move away from the more clearly defined open cellular convection cases towards the much more poorly defined and more disorganized clustered cumulus cases, they move towards a more central location in the $\mathcal{L}$-N_d joint distribution.

Overall, these highly granular cloud categories provide a clear picture of how different cloud types populate the $\mathcal{L}$-N_d space. The shape of the distribution of k-means clusters in this space (dot markers in Figure 3) matches well with the total joint PDF of $\mathcal{L}$ and N_d, which is unsurprising. However, there is not clear evidence of an inverted-V shape in the distribution of the k-means clusters. A limitation of this analysis is that the inverted-V pattern arises in conditional joint PDFs of $\mathcal{L}$ and N_d, and the locations of clusters in Figure 3 represent within-cluster means. Conditioning of $\mathcal{L}$ on N_d is only indirectly achieved here because each cluster tends to have N_d values in a specific range, so it is possible that an inverted-V pattern is less prominent here because conditioning on N_d is not done with the same regularity and fidelity as can be achieved with a conditional joint

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

histogram. We interpret this plot as evidence that while different cloud types tend to have unique $\mathcal{L}$ and N_d properties, the inverted-V pattern likely does not primarily arise from the distribution of cloud regimes in $\mathcal{L}$-N_d space.

4.2 The "inverted-V" within cloud regimes

In addition to examining the distribution of cloud morphology within $\mathcal{L}$-N_d space, we use the cloud clusters to investigate the $\mathcal{L}$-N_d distribution within different cloud regimes. Figure 4 contains a panel for each of the cloud types from Section 3.2 that shows the $\mathcal{L}$-N_d conditional PDF (CPDF) for the 25% of cases most similar (by Eq. 2) to their cluster center. We use only the top 25% to ensure we are examining cases that truly match the cloud type and avoid cases near the edge of the cluster that may be a poor match. Each $\mathcal{L}$ and N_d sample represents a value averaged over the pixels in an image chip with valid cloud property retrievals.

It is immediately noticeable that almost every cloud category shows a pattern resembling the inverted-V. For cloud categories that occupy mostly high- or mostly low-N_d parts of the space (refer to Figure 3), there are a limited number of samples in the most distant part of the space and the CPDFs become noisy there, but a partial inverted-V pattern is still evident. For example, the open cellular convection category (panel a) tends to have low N_d values, based on the previous section, and it largely occupies the ascending (left) leg and the peak of the inverted-V. There are samples of this cloud type with more moderate N_d values, however, and this region of panel a shows a clear decreasing trend in $\mathcal{L}$ with increasing N_d, with the CPDF eventually becoming noisy in the high-N_d region that has fewer samples. Likewise, closed cellular convection (panel b) mostly occupies the right part of this space where $\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}$ is negative, but the handful of closed cellular convection examples with $N_d < 30cm^{-3}$ still indicate an increasing slope. We suspect this is because most of this cloud regime is either not precipitating or only lightly precipitating, but this would require additional data to confirm. These examples are consistent with observations of ship tracks, that indicate a reduction in $\mathcal{L}$ in closed cell clouds polluted by ship emissions but an increase in $\mathcal{L}$ in open cell clouds polluted by ship emissions (Christensen and Stephens, 2011).

A notable exception to the inverted-V pattern is the bimodal distribution of $\mathcal{L}$ and N_d in the subtropical stratocumulus category. The downward sloping, high-N_d part of the inverted-V pattern is still there, along with a peak around 20-30 cm^{-3}, but there is also a low-N_d mode where $\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}$ is slightly negative at very low N_d's. This category mostly consists of very thin stratocumulus that tend to occur in the subtropics. Based on Figure 3, they mostly populate the high-N_d low-$\mathcal{L}$ part of the space, so the right part of the plot is more representative of how these clouds respond to changing N_d. We suspect the tendency for clouds in this category to be thin may also be the reason for the lack of a positive slope at very low N_d values, because thinner, non-precipitating, clouds will not show the precipitation suppression effect, but this would require further investigation with precipitation observations to verify.

Finally, the shape and location of the inverted-V pattern changes between different cloud types. For each of the panels in Figure 4 except e, we have estimated the slope $\left(\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}\right)$ for each branch of the V. We chose the N_d bin with the highest average $\mathcal{L}$ value as the cutoff between the ascending (red) and descending (blue) legs, and estimate the slope separately on each side. Some cloud types show relatively little $\mathcal{L}$ sensitivity to N_d, while some are much more sensitive. The stratus category, which is primarily composed of thick stratiform cloud cases, and the open cell category, have relatively weak slopes in the

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

descending branch for instance, while other categories show relatively strong responses. Additionally, the location of the V varies between cloud types. For most cloud types there is an obvious peak in $\mathcal{L}$ around $N_d = 20cm^{-3}$, but the storm track stratocumulus category has more of a plateau between $N_d = 10cm^{-3} - 70cm^{-3}$ and "sugar" has a more rounded shape than a V. This has implications for understanding potential changes in cloud properties in response to changes in aerosol loading, and studies focusing on aerosol impacts on a particular region and/or cloud type (e.g. subtropical stratocumulus decks) may not be generalizable to other cloud types. It also has implications for ACI estimates from ESMs because only recently have very high resolution cloud resolving ESMs (Caldwell et al., 2021; Donahue et al., 2024) begun to produce climate simulations with realistic marine boundary layer cloud morphologies, and even higher resolutions may still be necessary to fully capture marine boundary layer stratocumulus (Bogenschutz et al., 2021).

Goren et al. (2025) and Goren et al. (2026) have also recently studied the relationship between cloud morphology and $\mathcal{L}$ susceptibility to N_d, and have similar, though not identical, findings. Goren et al. (2025) showed that the high-N_d portion of the inverted-V pattern is dominated by closed cells while the low-N_d portion is dominated by open cells. Our results confirm this and show that other cloud types occupy specific regions in $\mathcal{L}$-N_d space as well. Goren et al. (2025) also indicate that the inverted-V seen in past studies is largely due to aggregating multiple cloud regimes in the analysis though, which Figure 4 indicates is only partially true: closed cell cases have a weak positive $\mathcal{L}$-N_d relationship at low N_d, open cell cases have a clear negative branch for high-N_d, and most cloud regimes show evidence of the inverted-V pattern. We also note that Goren et al. (2025) demonstrates that co-variability between PBL height and droplet number may explain the negative portion of the $\mathcal{L}$-N_d relationship. We do not directly investigate this here, but supplement S1 shows $\mathcal{L}$-N_d distributions for the more granular clusters produced by k-means, and regional variability of $\mathcal{L}$-N_d properties of clouds within each of the super-clusters here is evident. Goren et al. (2026) investigate the relationship in terms of cloud morphology defined by the distribution of cloud optical thicknesses within a scene and show almost all morphologies defined this way have a negative $\mathcal{L}$-N_d slope. While Figure 4 shows some evidence of a positive-slope branch of the inverted-V in most cloud regimes, Figure 3 shows that all cloud types except open cellular convection primarily occur in the negative-slope portion of the V.

Overall, we interpret Figure 4 as further evidence that the inverted-V pattern does not occur as a result of the distribution of different cloud types in $\mathcal{L}$-N_d space because individual cloud regimes each show evidence of an inverted-V pattern. That said, different cloud types dominate different locations in the phase space and there are clear differences in the exact slope and shape of the inverted-V between different cloud regimes, which implies that when assessing expected responses of cloud radiative forcing to changing aerosol loading, taking cloud morphology into account is important.

4.3 Explaining $\mathcal{L}$ variance as a function of cloud morphology

This section analyzes the $\mathcal{L}$-N_d relationship while accounting for cloud type by leveraging the continuous vector embedding representations of cloud morphology discussed in Section 3.3 rather than relying on discrete representations from clustering or hand labeling. Small feed-forward multi-layer perceptron type neural networks (Bishop, 2006) are used as generic non-linear fitting tools to find continuous and differentiable mappings between cloud organizational structure and $\mathcal{L}$, and then analyzed to separate contributions to the inverted-V pattern from N_d and cloud type. The neural networks are trained to map between

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

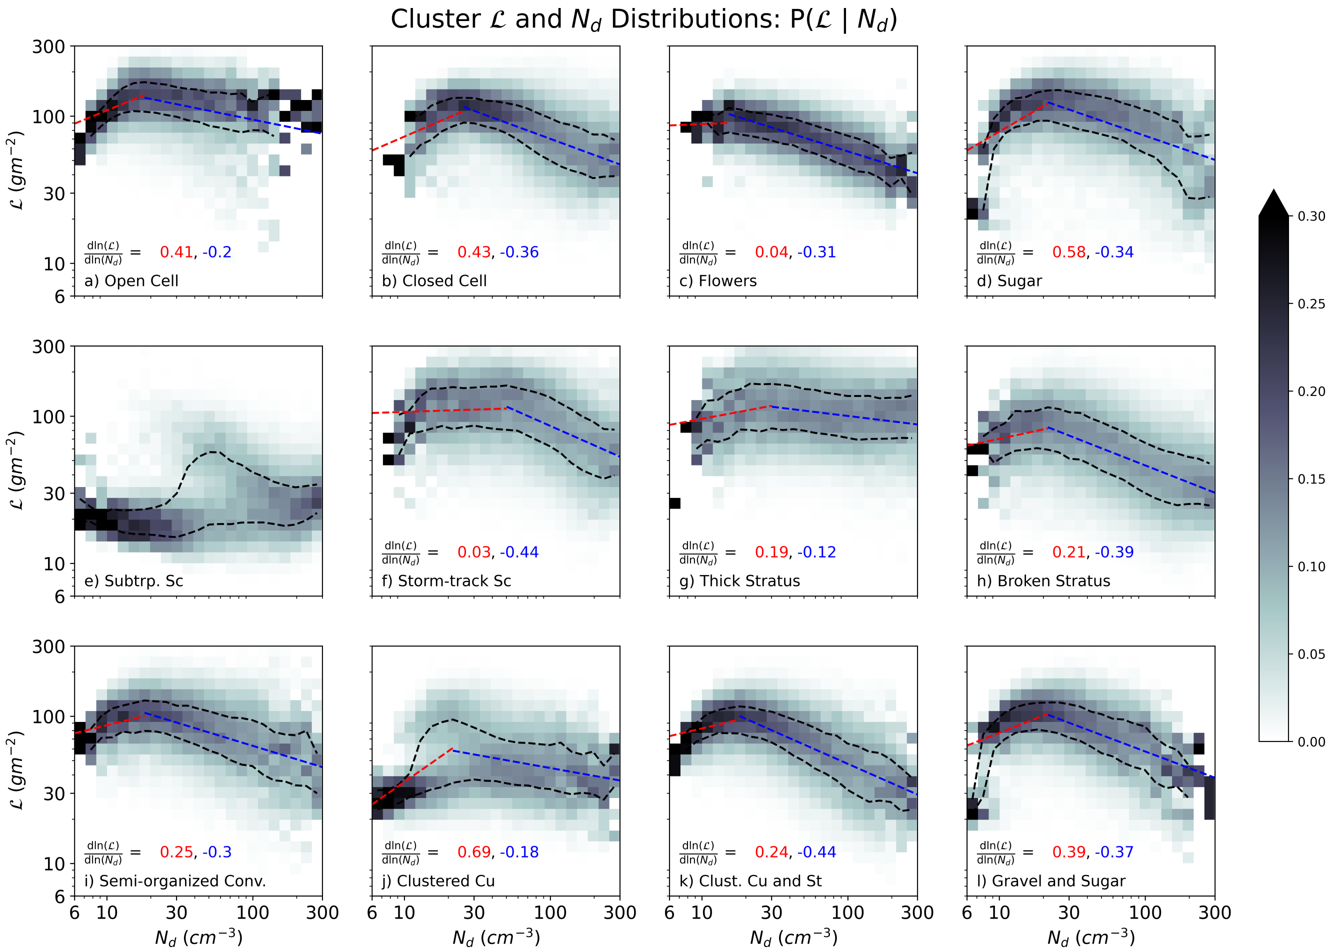

Figure 4. The $\mathcal{L}$ - N_d relationship within different cloud morphology clusters shown as the conditional probability distribution $P(\mathcal{L}|N_d)$. Shading indicates the probability density of samples with respect to $\mathcal{L}$ for each N_d bin. Red and blue dashed lines show fits to the slope $\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}$ for each side of the "inverted-V". Dashed black lines represent the inter-quartile range of $\mathcal{L}$ values for each N_d bin.

various sets of inputs to $\ln(\mathcal{L})$. The predictors explored are: $\ln(N_d)$, $\mathbf{x}$, $\widetilde{\mathbf{x}}$, λ, and ϕ. Here, λ and ϕ represent the longitude and latitude respectively, and are included because Goren et al. (2025) identified a strong covariance between location within an ocean basin and the $\mathcal{L}$-N_d relationship. $\mathbf{x}$ represents the 1024-value vector embeddings of cloud morphology from the self supervised algorithm. Because these are very long vectors, and might over-parameterize cloud morphology, potentially leading to over-fitting or drowning out the signal from other input variables, we also use $\widetilde{\mathbf{x}}$, which is the first 32 principal components (Bishop, 2006) of the morphology vector embedding dataset. They describe 92.5% of the variance in the vector embedding dataset. Because the scenes analyzed here are heavily filtered to only include cases with a sufficient amount of low liquid clouds (Section 2.2) and the self-supervised model was trained on all MODIS cloud imagery globally for a year, the diversity

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

in the imagery used in this study is much lower than the diversity in the imagery dataset that the self-supervised algorithm in Geiss et al. (2024) was trained on. Thus, only a small number of principal components are needed to capture most of the variance. Each of the inputs to the neural network are scaled to have a mean of zero and unit variance. This is done individually for $\ln(N_d)$, λ, and ϕ. For $\mathbf{x}$, the direction of the vector is more representative of image content than the magnitude (this is why the cosine similarity metric is used to compare vector embeddings), so these are scaled by dividing by the standard deviation of all vector components over all samples, which preserves the orientation of the vectors (as opposed to scaling individual vector components). Ultimately, nine versions of the model are trained, one for each unique input (λ and ϕ are always used together), and one for each unique pair of inputs except $\mathbf{x}$ and $\widetilde{\mathbf{x}}$, which contain redundant information. For each trained model, we analyze the fraction of $\mathcal{L}$ variance explained, defined:

$$1 - \frac{\sum_i \left(\ln(\mathcal{L}_i) - \widehat{\ln(\mathcal{L}_\mathrm{i})}\right)^2}{\sum_i \left(\ln(\mathcal{L}_i) - \overline{\ln(\mathcal{L})}\right)^2} \tag{4}$$

where the overline represents a mean and the hat indicates predicted values from the neural network.

The neural networks are fully connected feed-forward models with 3 hidden layers that use sigmoid linear unit activations, also called “SiLU” or “swish” (Hendrycks and Gimpel, 2016). They are trained to predict the natural logarithm of the liquid water path which has been linearly scaled to a range of (-1,1) so the neural network can apply a tanh activation to its output. The neural networks are trained using mean squared error loss and the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 10^{-3} and otherwise using Keras (Chollet et al., 2015) default settings. Each neural network is trained for a minimum of 10 and a maximum of 100 epochs. After 3 consecutive epochs without an improvement in validation score the model weights from the best performing epoch are saved. This early stopping based on validation set performance helps ensure the models are not overfit, though overfitting is unlikely due to the fact that the models are small compared to the number of samples and the diversity in the dataset (1.4×10^5 trainable parameters for the largest model that ingests $\mathbf{x}$ or 8.4×10^3 for the models that take $\widetilde{\mathbf{x}}$ as an input compared to approximately 3.4×10^6 total samples).

To assess uncertainty in estimates of LWP variance explained by the neural networks, we train multiple copies of each using a k-folds cross validation strategy (Kohavi, 1995). The full dataset is composed of 10 full years of data. The last two years, 2021 and 2022, are used for testing of all models with 2013-2020 used for training and validation. For each k-folds training, we hold out one year of data from 2013-2020 to use for validation and train on the remainder, resulting in 8 splits. This process approximates the noise in the estimates of percent variance explained by the each model configuration that might be introduced by selection of the validation and training sets, model weight initialization, or the stochastic neural network training process.

Table 1 shows the fractional variance explained by each neural network using different combinations of predictor variables. The table entries represent the mean and a 95% confidence interval based on the 8 iterations of k-folds. Using $\ln(N_d)$ alone as a predictor only explains 4.33% of the variance in $\ln(\mathcal{L})$. It can be seen by examination of Figure 4, and throughout the literature on this topic, that there is clearly a robust non-linear relationship between these two variables, but it can also be seen in these figures that there is substantial spread in $\mathcal{L}$ when plotted as a function of N_d, which is the reason for the low predictability.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

Table 1. Percent variance explained by an MLP trained to predict $\ln(\mathcal{L})$ from the indicated predictor variables. $\mathbf{x}$ is a 1024-length vector representing cloud morphology and $\widetilde{\mathbf{x}}$ is a 32-length vector. The table indicates the mean value and a 95% confidence interval based on the 8-fold cross validation.

First predictor:	Second predictor:		
	None	$\ln(N_d)$	λ, ϕ
$\ln(N_d)$	4.33 ± 0.07	-	-
λ, ϕ	8.43 ± 0.23	14.46 ± 0.41	-
$\widetilde{\mathbf{x}}$	69.47 ± 0.09	73.68 ± 0.13	71.33 ± 0.15
$\mathbf{x}$	75.21 ± 0.26	79.22 ± 0.31	76.61 ± 0.22

The latitude and longitude, representing the location within the ENA region, are stronger predictors for $\mathcal{L}$ than N_d, explaining 8.43% of the variance. This is consistent with Goren et al. (2025) who identify location, and particularly longitude, as a strong confounding variable that is correlated with both N_d and $\mathcal{L}$. Their study regions are primarily subtropical eastern boundaries of the world's ocean basins – areas dominated by marine stratus and where longitude represents distance from the coast and thus will be a good predictor of SST and marine boundary layer properties. Using $\widetilde{\mathbf{x}}$ or $\mathbf{x}$ as predictors explain significantly more variance in $\ln(\mathcal{L})$, 69.47% and 75.21% respectively. This means that the majority of $\mathcal{L}$ variance between scenes can be explained simply as a function of cloud morphology. While past studies have investigated how cloud type might be correlated with liquid water path, and quantified typical $\mathcal{L}$ within specific cloud regimes, the use of a *continuous* representation of cloud morphology as a predictor for $\mathcal{L}$ here goes a step beyond this because it is not limited to a small number of discrete cloud regimes and is less influenced by human biases.

The models trained with different combinations of input variables provide further insight. The model trained with both $\widetilde{\mathbf{x}}$ and location information as inputs explains only an additional 1.86% of $\ln(\mathcal{L})$ variance than the model trained with $\widetilde{\mathbf{x}}$ as the only input, while the model trained with only location information as the input explained more than 8% of the $\ln(\mathcal{L})$ variance. This indicates that the variance in $\ln(\mathcal{L})$ that can be explained by λ and ϕ is primarily due to covariance of cloud type and location. Meanwhile, when used as the sole predictor variable, $\ln(N_d)$ describes only 4.33% of the variance in $\ln(\mathcal{L})$, but when combined with $\widetilde{\mathbf{x}}$ as an extra predictor, the variance explained by the model increases by 4.21%. A similar result can be seen using the full vector embeddings ($\mathbf{x}$) as inputs. This indicates that the $\mathcal{L}$ variance explained by N_d is largely independent of cloud type. Overall, the results in this section indicate that while morphology can explain the majority of the variance in cloud $\mathcal{L}$, the portion explained by N_d is mostly independent of cloud type.

4.4 Using model interrogation to account for cloud type

To the extent that the neural networks in the previous section have accurately learned the explainable portion of the relationship between N_d, $\widetilde{\mathbf{x}}$, and $\mathcal{L}$, we can interrogate them to learn about this relationship in ways that are difficult with raw data. There

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

are two approaches presented here: 1) assessing the residual of predictions of the neural networks from the previous section compared to ground truth, and 2) differentiation of the neural networks.

4.4.1 Neural network residuals

Figure 5a is the CPDF of $\ln(\mathcal{L})$ data conditioned on $\ln(N_d)$ in the neural network testing dataset, and demonstrates the inverted-V pattern. Panels b-d in this figure, show the residual component of $\ln(\mathcal{L})$ that cannot be predicted by three of the neural networks from the previous section. Specifically, they show $P((\ln(\mathcal{L}) - \widehat{\ln(\mathcal{L})})|\ln(N_d))$ from the neural networks that predict $\ln(\mathcal{L})$ using $\ln(N_d)$, $\widetilde{\mathbf{x}}$, or both as inputs, with the input variables for the neural networks shown in the panel titles. The networks trained on the first k-fold split were used to produce this figure. Panel b, shows the residual CPDF when using a neural network trained to predict $\ln(\mathcal{L})$ from $\ln(N_d)$ only. The inverted-V pattern is gone, but a great deal of spread in $\ln(\mathcal{L}) - \widehat{\ln(\mathcal{L})}$ remains. Panel c shows the component of $\ln(\mathcal{L})$ that cannot be explained by the neural network model trained to only take $\widetilde{\mathbf{x}}$ as input. In this case, there is less spread, but the inverted-V pattern remains. These can be compared to panel d, which shows a similar plot for the neural network model that takes both $\ln(N_d)$ and $\widetilde{\mathbf{x}}$ as inputs. While there is still a clear inverted-V pattern in panel c, in panel d this pattern is missing because it can be predicted by the neural network. This is an indication that the inverted-V pattern is not explainable as a function of cloud type alone, but is explainable when N_d data are provided.

4.4.2 Model differentiation

A valuable feature of neural networks is that they are continuous and differentiable approximate representations of the underlying training data, and modern neural network libraries (Chollet et al., 2015) allow automatic differentiation of neural network models. The derivative $\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}$ is important to understanding climate responses to changes in aerosol loading, and it is a contributor to the effective radiative forcing due to aerosol cloud interaction. Estimates of this value from observations typically involve a linear, or piecewise-linear fit to data, and because of the non-linearity of the relationship, some amount of binning with respect to $\mathcal{L}$ or another cloud variable might also be used. However, these linear fits aggregate data across multiple cloud types. Goren et al. (2025) showed, as we have shown above, that cloud type is a strong confounding factor for this relationship, which means that such fits are at least partially indicative of the distribution of cloud types in $\mathcal{L}$-N_d space instead of the response of individual clouds to changing N_d. Here, we use the trained neural network as a representation of the $\mathcal{L}$-N_d-cloud morphology relationship, and interrogate the neural network as a way to control for cloud type as a confounder, by calculating the derivative of the network trained to map $\ln(N_d)$ and $\widetilde{\mathbf{x}}$ to $\ln(\mathcal{L})$ with respect to $\ln(N_d)$ using automatic differentiation. Specifically, for each sample in the dataset, we compute $\left.\frac{\partial\widehat{\ln(\mathcal{L})}}{\partial\ln(N_d)}\right|_{\widetilde{x}}$, where the hat indicates that we are differentiating the value of $\mathcal{L}$ predicted by the neural network and not the true value. For each sample in the dataset, this gives an estimate of how much the liquid water path is expected to change in response to changing N_d when cloud type is not allowed to change.

We use this approach to compute the $\mathcal{L}$-N_d derivative for each sample in the dataset (independent of the cloud regime labels) and Figure 6 shows these values binned and averaged in $\mathcal{L}$-N_d space. Stippling in the figure indicates bins where the slope is non-zero with 99% confidence (using a t-test). Overall, we see a pattern that is consistent with the inverted-V pattern, with

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

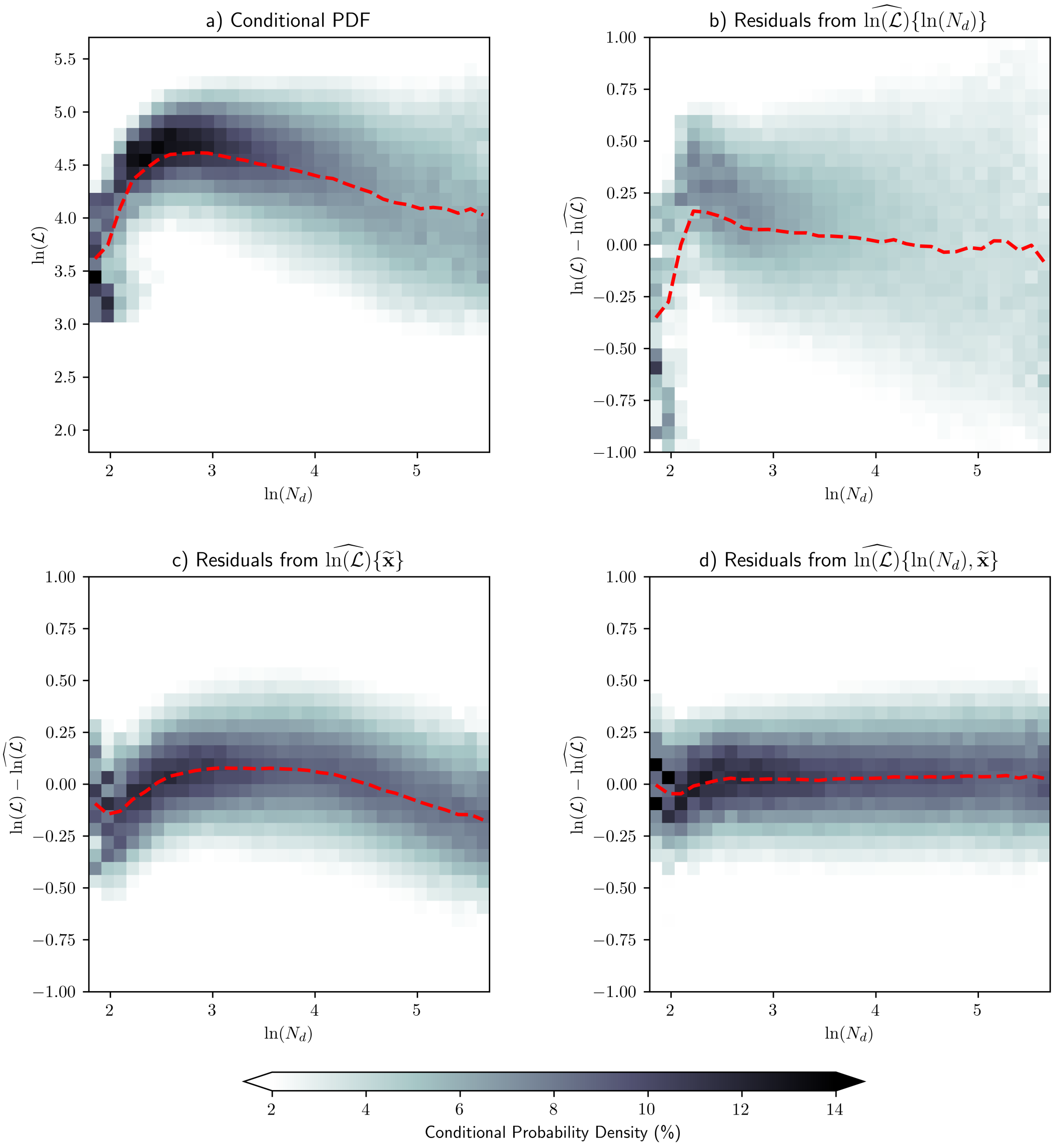

Figure 5. Experiments showing the residual $\ln(\mathcal{L})$ that cannot be predicted by different neural network configurations. Panel a) is meant for reference and shows the CPDF of $\ln(\mathcal{L})$ and $\ln(N_d)$ to demonstrate the inverted-V pattern. The red dashed line in each panel shows the mean $\ln(\mathcal{L})$ (or $\ln(\mathcal{L}) - \widehat{\ln(\mathcal{L})}$) value for each $\ln(N_d)$ bin. Hats represent variables that are predicted by the neural networks. Panel b) shows the residual after predicting $\ln(\mathcal{L})$ with a neural network that takes only $\ln(N_d)$ as an input. Panel c) is similar but for a neural network that takes only $\widetilde{\mathbf{x}}$ (a 32-value vector embedding representing cloud morphology) as an input. Panel d) is similar but for a neural network that takes both $\ln(N_d)$ and $\widetilde{\mathbf{x}}$ as inputs.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

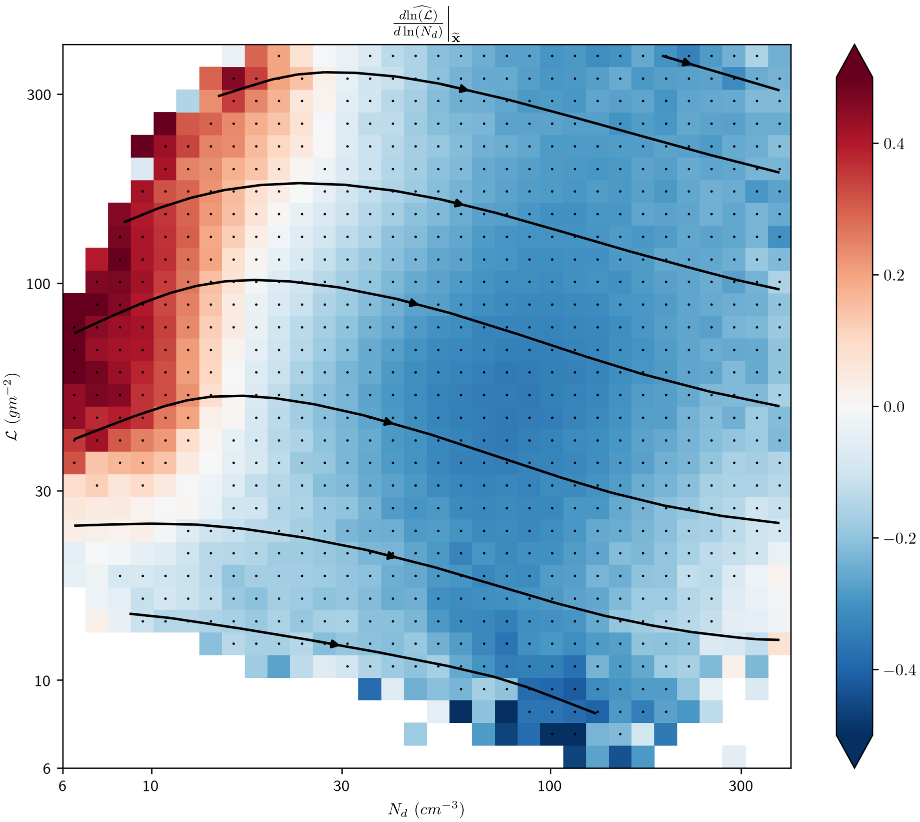

Figure 6. Derivatives of $\frac{d\ln(\widehat{\mathcal{L}})}{d\ln(N_d)}\Big|_{\widetilde{x}}$ aggregated across the cloud dataset (including all cloud morphologies) and binned with respect to $\mathcal{L}$ and N_d. This computed by training a neural network to predict $\ln(\mathcal{L})$ from $\ln(N_d)$ and a 32-value vector embedding representing cloud morphology ($\widetilde{x}$), differentiating the neural network at the location of individual samples, and aggregating the results. Streamlines indicate the same slopes shown by the color map and give a visual indication of how we would expect clouds to move in this phase space if N_d were increased without allowing the cloud morphology to change (with the caveat that this is only reliable for small deviations in N_d). Stippling represents bins that contain a non-zero slope with 99% confidence.

LWP increasing with respect to N_d for low-N_d clouds and decreasing at higher N_d values. This indicates that this relationship is robust even when the impact of different cloud morphologies is controlled for. Figure 6 also shows streamlines which follow the underlying slope data. These streamlines give a visual indication of how we would expect clouds to move in this phase space if N_d were increased without allowing the cloud morphology to change (with the caveat that this is only reliable for small deviations in N_d). Again, the inverted-V pattern is robust, indicating that it arises largely independently of cloud type.

5 Discussion

The above analysis provides several insights into the $\mathcal{L}$-N_d relationship and how it varies across cloud regimes:

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

1. **Different cloud morphologies tend to occupy different regions within the $\mathcal{L}$-N_d space.** This has been documented before with a supervised machine learning approach using three cloud categories (Goren et al., 2025), but this study further cements this concept by showing it is true across a larger range of different morphological regimes identified using a self supervised machine learning scheme, and is even true when we consider a very high number of ML-identified cloud regimes (64).

2. **The "inverted-V" shape likely does not appear as a result of (1).** Firstly, the distribution of different cloud types in the $\mathcal{L}$-N_d space does not show clear evidence of an inverted-V pattern (Figure 3). Secondly, when the $\mathcal{L}$-N_d joint CPDF is plotted within each cloud regime (Figure 4), we see that the inverted-V manifests in almost all cloud regimes, though the ascending and descending branches can have different locations and slopes depending on cloud type, and both branches/legs are not well represented in all cloud regimes (clouds that are typically non-precipitating show only weak evidence of the upward, low-N_d, leg, for instance).

3. **Cloud morphology can explain much more of the variance in $\mathcal{L}$ than N_d, but the portion of $\mathcal{L}$ variance explained by these two factors is largely independent.** We demonstrated that a simple neural network regressor can explain 69.47% of the variance in $\ln(\mathcal{L})$ across different cloud scenes using only a 32-length vector encoding cloud morphology as an input. Meanwhile, using N_d as a predictor only explains 4.33% of the $\ln(\mathcal{L})$ variance. However, using both predictors results in 73.68% variance explained, indicating that most of the $\mathcal{L}$ variance explained by N_d cannot be explained by cloud morphology information.

4. **When ML models are used to control for the impacts of different cloud morphologies, the inverted-V relationship between N_d and LWP remains robust.** Interrogation of trained neural network regressors from section 4.3 allows us to estimate the value of $d\ln(\mathcal{L})/d\ln(N_d)$ while holding cloud morphology constant at the level of individual image samples in the dataset, and the result shows a strong inverted V pattern. Furthermore, neural network models trained to predict $\ln(\mathcal{L})$ from cloud morphology alone leave a residual inverted V pattern that is not explained, while models that ingest both morphology and N_d information can predict this pattern.

A major limitation of this study, and others on this topic, is the dependence on passive satellite $\mathcal{L}$ and N_d retrievals. The potential sources of uncertainty in these retrievals are discussed in detail in Section 2.2. In short, assumptions about cloud homogeneity, phase, and adiabaticity, viewing and solar zenith angle changes, and dependence of both retrievals on τ_c and r_e are sources of uncertainty. While strict filtering of the data can mitigate many of these issues, our results should still be interpreted with this in mind. One potential unique source of uncertainty in this study, is the relatively high granularity of the cloud regimes studied. Some sources of retrieval uncertainty are much more common in specific cloud regimes. For instance, the "sugar" category has a much higher fraction of partly cloudy pixels then the "closed cellular convection" category.

While satellite retrievals have unique uncertainties, there is also evidence of the inverted-V pattern from ground-based remote sensing (Tian et al., 2025; Varble et al., 2023; Tang et al., 2023; Wu et al., 2020) though it is less clear and can have different slopes. These approaches do away with the adiabaticity (f_{ad}) assumption required for the passive satellite retrieval

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

of N_d, but introduce their own sources of uncertainty. It has also been demonstrated in modeling studies (Mülmenstädt et al., 2024; Kokkola et al., 2025). In light of these limitations, future work in this area can focus on integrating observations from ground based instruments at the ARM ENA site to verify these results. This site houses a microwave radiometer and multi-filter rotating shadow band radiometer that can be used to perform ground-based $\mathcal{L}$ and N_d retrievals (Zhang et al., 2020; Riihimaki et al., 2021). Additionally, Tian et al. (2025) demonstrated that at least open- and closed-cellular convection cloud types, can be reliably classified using the Ka-band zenith radar at the site. Applying a self-supervised scheme like Geiss et al. (2024) to radar and/or geostationary imager data from the site could also allow for creation of a higher time-resolution cloud type dataset rather than relying on MODIS overpasses. Use of geostationary data will also enable Lagrangian tracking of clouds and can enable study of transitions between cloud regimes and the $\mathcal{L}$ response to N_d as a function of time (Christensen et al., 2023, 2024).

There are several other potential avenues of future investigation. Firstly, the self-supervised morphology classification scheme can be trivially modified to evaluate representation of cloud regimes and their properties and responses to cloud controlling factors in models, at least those with sufficient fidelity to resolve diverse cloud regimes. This can be used to assess potential model shortcomings in representing marine stratiform clouds and their responses to different forcings. Secondly, the variations in the inverted-V pattern shown in section 4.2 warrant further investigation. These differences are likely a function of meteorological factors that can both determine the type of cloud regime present but also how clouds interact with the surrounding atmosphere (e.g. due to humidity, stability etc.), and potentially cloud properties that were not evaluated here, such as vertical structure and precipitation. Finally, the method used in Section 4.4.2 for estimating $\left.\widehat{\frac{\mathrm{d}\ln(\mathcal{L})}{\mathrm{d}\ln(N_d)}}\right|_{\tilde{x}}$ directly for individual samples using a trained neural network is novel, and has potential for much broader use in climate science. Estimates of radiative forcing due to changes in key climate variables using derivatives estimated from observations and models are common in climate science (Held and Soden, 2000; Zelinka et al., 2012; Hartmann, 2015; Bellouin et al., 2020). This method offers a unique way to control for confounding factors, handle non-linearity, and to estimate individual components of chain-rule expansions of key relationships (e.g. Bellouin et al. (2020)). The main drawbacks of this approach are that it requires training a neural network and its accuracy will be limited by the accuracy of the neural network, and that uncertainty estimates are more difficult than with classical linear fits. Nonetheless, there is significant potential to deploy this method to better understand the radiative effects of clouds and aerosols, particularly in the context of different cloud morphologies.

Code and data availability.

We have made the code, intermediate data products (not including the MODIS granule files), and neural networks produced for the analysis in this work publicly available in a Zenodo repository (Geiss, 2026): https://zenodo.org/records/18237664. RESTRICTED ZENODO ACCESS DURING REVIEW: https://zenodo.org/records/18237664?preview=1&token=eyJhbGciOiJIUzUxMiJ9.eyJpZCI6ImIxOTQ2OWMwLTRlYzAtNDMxYi1iNWNhLWExM2IyOTlmNWRiNyIsImRhdGEiOnt9LCJyYW5kb20iOiJhNTU3ZTk1NDMzOWUxNmFhNjczNTEzMzRlN2UwOTVlNyJ9.SiSRScGHiaUzl1Gza3xrNMAAwijwLHWNWdMN2wvc5pCIX8PvosQrHuFfp9Yryfcu2YFt05XEREBNcjGmmM_egw

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

The MODIS datasets are too large to include, but are publicly available via LAADS DAAC (see below), and the Zenodo repository includes lists of file names of the MODIS granules we used. The Terra MODIS L1 reflectance imagery have DOI: 10.5067/MODIS/MOD02HKM.061 and can be downloaded from: https://ladsweb.modaps.eosdis.nasa.gov/missions-and-measurements/products/MOD02HKM (Last access November 6^{th} 2025). The Terra MODIS L2 cloud data used here have DOI: 10.5067/MODIS/MOD06_L2.061 and can be downloaded from: https://ladsweb.modaps.eosdis.nasa.gov/missions-and-measurements/products/MOD06_L2#overview (Last access November 6^{th} 2025).

A trained copy of the self-supervised CNN used to classify reflectance imagery can be found in the Zenodo repository from Geiss et al. (2024): https://zenodo.org/records/17137530

Author contributions. AG performed analysis and wrote the manuscript, JM conceived study and performed analysis, HX aided with expert clustering, all authors provided input on methods and results and edited the manuscript.

Competing interests. Matthew Christensen is a member of the ACP editorial board. Otherwise the authors declare no competing interests.

Acknowledgements. This research was a component of the Integrated Cloud, Land-surface, and Aerosol System Study (ICLASS) supported by the U.S. Department of Energy Office of Science Biological and Environmental Research as part of the Atmospheric Systems Research (ASR) Program under project KP1701010/57131. Pacific Northwest National Laboratory is operated by Battelle for the U.S. Department of Energy under Contract DE-AC05-76RLO1830.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

References

Ackerman, A. S., Kirkpatrick, M. P., Stevens, D. E., and Toon, O. B.: The impact of humidity above stratiform clouds on indirect aerosol climate forcing, Nature, 432, 1014–1017, 2004.

Albrecht, B. A.: Aerosols, Cloud Microphysics, and Fractional Cloudiness, Science, 245, 1227–1230, https://doi.org/10.1126/science.245.4923.1227, 1989.

Arola, A., Lipponen, A., Kolmonen, P., Virtanen, T. H., Bellouin, N., Grosvenor, D. P., Gryspeerdt, E., Quaas, J., and Kokkola, H.: Aerosol effects on clouds are concealed by natural cloud heterogeneity and satellite retrieval errors, Nature Communications, 13, 7357, https://doi.org/10.1038/s41467-022-34948-5, 2022.

Bellouin, N., Quaas, J., Gryspeerdt, E., Kinne, S., Stier, P., Watson-Parris, D., Boucher, O., Carslaw, K. S., Christensen, M., Daniau, A.-L., et al.: Bounding global aerosol radiative forcing of climate change, Reviews of Geophysics, 58, e2019RG000 660, 2020.

Bender, F.-M., Frey, L., McCoy, D. T., Grosvenor, D. P., and Mohrmann, J. K.: Assessment of aerosol–cloud–radiation correlations in satellite observations, climate models and reanalysis, Climate Dynamics, 52, 4371–4392, 2019.

Bishop, C.: Pattern Recognition and Machine Learning, Information Science and Statistics, Springer, ISBN 9780387310732, https://books.google.com/books?id=qWPwnQEACAAJ, 2006.

Bogenschutz, P. A., Yamaguchi, T., and Lee, H.-H.: The Energy Exascale Earth System Model simulations With high vertical resolution in the lower troposphere, Journal of Advances in Modeling Earth Systems, 13, e2020MS002 239, 2021.

Boucher, O., Randall, D., Artaxo, P., Bretherton, C., Feingold, G., Forster, P., Kerminen, V.-M., Kondo, Y., Liao, H., Lohmann, U., et al.: Clouds and aerosols, in: Climate change 2013: The physical science basis. Contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change, pp. 571–657, Cambridge University Press, 2013.

Bradley, P. S., Fayyad, U. M., et al.: Refining initial points for k-means clustering., in: ICML, vol. 98, pp. 91–99, Citeseer, 1998.

Bretherton, C., Blossey, P. N., and Uchida, J.: Cloud droplet sedimentation, entrainment efficiency, and subtropical stratocumulus albedo, Geophysical research letters, 34, https://doi.org/10.1029/2006GL027648, 2007.

Caldwell, P. M., Terai, C. R., Hillman, B., Keen, N. D., Bogenschutz, P., Lin, W., Beydoun, H., Taylor, M., Bertagna, L., Bradley, A., et al.: Convection-permitting simulations with the E3SM global atmosphere model, Journal of Advances in Modeling Earth Systems, 13, e2021MS002 544, 2021.

Celebi, M. E., Kingravi, H. A., and Vela, P. A.: A comparative study of efficient initialization methods for the k-means clustering algorithm, Expert Systems with Applications, 40, 200–210, https://doi.org/10.1016/j.eswa.2012.07.021, 2013.

Chen, T., Kornblith, S., Norouzi, M., and Hinton, G.: A simple framework for contrastive learning of visual representations, in: International conference on machine learning, pp. 1597–1607, PmLR, 2020.

Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., and Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets, Advances in neural information processing systems, 29, 2016.

Chen, Y.-C., Christensen, M. W., Stephens, G. L., and Seinfeld, J. H.: Satellite-based estimate of global aerosol–cloud radiative forcing by marine warm clouds, Nature Geoscience, 7, 643–646, 2014.

Chollet, F. et al.: Keras, https://github.com/fchollet/keras, 2015.

Christensen, M. W. and Stephens, G. L.: Microphysical and macrophysical responses of marine stratocumulus polluted by underlying ships: Evidence of cloud deepening, Journal of Geophysical Research: Atmospheres, 116, https://doi.org/https://doi.org/10.1029/2010JD014638, 2011.

Christensen, M. W., Ma, P.-L., Wu, P., Varble, A. C., Mülmenstädt, J., and Fast, J. D.: Evaluation of aerosol–cloud interactions in E3SM using a Lagrangian framework, Atmospheric Chemistry and Physics, 23, 2789–2812, https://doi.org/10.5194/acp-23-2789-2023, 2023.

Christensen, M. W., Wu, P., Varble, A. C., Xiao, H., and Fast, J. D.: Aerosol-induced closure of marine cloud cells: enhanced effects in the presence of precipitation, Atmospheric Chemistry and Physics, 24, 6455–6476, https://doi.org/10.5194/acp-24-6455-2024, 2024.

Denby, L.: Discovering the importance of mesoscale cloud organization through unsupervised classification, Geophysical Research Letters, 47, e2019GL085 190, 2020.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L.: Imagenet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255, Ieee, 2009.

Donahue, A. S., Caldwell, P. M., Bertagna, L., Beydoun, H., Bogenschutz, P. A., Bradley, A., Clevenger, T. C., Foucar, J., Golaz, C., Guba, O., et al.: To exascale and beyond—The Simple Cloud-Resolving E3SM Atmosphere Model (SCREAM), a performance portable global atmosphere model for cloud-resolving scales, Journal of Advances in Modeling Earth Systems, 16, e2024MS004 314, 2024.

Fan, J., Wang, Y., Rosenfeld, D., and Liu, X.: Review of aerosol–cloud interactions: Mechanisms, significance, and challenges, Journal of the Atmospheric Sciences, 73, 4221–4252, 2016.

Forster, P., Storelvmo, T., Armour, K., Collins, W., Dufresne, J.-L., Frame, D., Lunt, D., Mauritsen, T., Palmer, M., Watanabe, M., Wild, M., and Zhang, H.: Chapter 7: The Earth's energy budget, climate feedbacks, and climate sensitivity, IPCC AR6, https://doi.org/10.25455/wgtn.16869671, 2021.

Geiss, A.: Self Supervised Cloud Classification Data Supplement (Version 1) [Data set], https://doi.org/10.5281/zenodo.7823778, 2023.

Geiss, A.: Code and Data Supplement for: A machine learning approach to understanding the relationship between cloud liquid water path and droplet number concentration across North Atlantic cloud regimes., https://doi.org/10.5281/zenodo.18237664, 2026.

Geiss, A., Christensen, M. W., Varble, A. C., Yuan, T., and Song, H.: Self-Supervised Cloud Classification, Artificial Intelligence for the Earth Systems, 3, e230 036, https://doi.org/10.1175/AIES-D-23-0036.1, 2024.

Glassmeier, F., Hoffmann, F., Johnson, J. S., Yamaguchi, T., Carslaw, K. S., and Feingold, G.: Aerosol-cloud-climate cooling overestimated by ship-track data, Science, 371, 485–489, https://doi.org/10.1126/science.abd3980, 2021.

Goodfellow, I., Bengio, Y., and Courville, A.: Deep Learning, MIT Press, http://www.deeplearningbook.org, 2016.

Goren, T., Choudhury, G., Kretzschmar, J., and McCoy, I.: Co-variability drives the inverted-V sensitivity between liquid water path and droplet concentrations, Atmospheric Chemistry and Physics, 25, 3413–3423, 2025.

Goren, T., Choudhury, G., and Feingold, G.: Beyond discrete stratocumulus regimes: a ternary continuum of morphology reveals within-regime variability in cloud susceptibilities, Atmospheric Chemistry and Physics, 26, 7193–7206, https://doi.org/10.5194/acp-26-7193-2026, 2026.

Grosvenor, D. P., Field, P. R., Hill, A. A., and Shipway, B. J.: The relative importance of macrophysical and cloud albedo changes for aerosol-induced radiative effects in closed-cell stratocumulus: insight from the modelling of a case study, Atmospheric Chemistry and Physics, 17, 5155–5183, 2017.

Grosvenor, D. P., Sourdeval, O., Zuidema, P., Ackerman, A., Alexandrov, M. D., Bennartz, R., Boers, R., Cairns, B., Chiu, J. C., Christensen, M., et al.: Remote sensing of droplet number concentration in warm clouds: A review of the current state of knowledge and perspectives, Reviews of Geophysics, 56, 409–453, 2018.

Gryspeerdt, E., Stier, P., and Partridge, D.: Satellite observations of cloud regime development: the role of aerosol processes, Atmospheric Chemistry and Physics, 14, 1141–1158, 2014.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

Gryspeerdt, E., Goren, T., Sourdeval, O., Quaas, J., Mülmenstädt, J., Dipu, S., Unglaub, C., Gettelman, A., and Christensen, M.: Constraining the aerosol influence on cloud liquid water path, Atmospheric Chemistry and Physics, 19, 5331–5347, https://doi.org/10.5194/acp-19-5331-2019, 2019.

Han, Q., Rossow, W. B., Zeng, J., and Welch, R.: Three different behaviors of liquid water path of water clouds in aerosol–cloud interactions, Journal of the atmospheric sciences, 59, 726–735, 2002.

Hartigan, J. A.: Clustering algorithms, John Wiley & Sons, Inc., 1975.

Hartmann, D. L.: Global physical climatology, vol. 103, Newnes, 2015.

Held, I. M. and Soden, B. J.: Water vapor feedback and global warming, Annual review of energy and the environment, 25, 441–475, 2000.

Hendrycks, D. and Gimpel, K.: Gaussian error linear units (gelus), arXiv preprint arXiv:1606.08415, https://doi.org/https://doi.org/10.48550/arXiv.1606.08415, 2016.

Houze, R.: Cloud Dynamics, International Geophysics, Elsevier Science, ISBN 9780123742667, https://books.google.com/books?id=aajnoAEACAAJ, 2014.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, https://doi.org/10.48550/ARXIV.1412.6980, 2014.

Kohavi, R.: A study of cross-validation and bootstrap for accuracy estimation and model selection, Ijcai, 14, 1137–1145, 1995.

Kokkola, H., Tonttila, J., Calderón, S. M., Romakkaniemi, S., Lipponen, A., Peräkorpi, A., Mielonen, T., Gryspeerdt, E., Virtanen, T. H., Kolmonen, P., and Arola, A.: Model analysis of biases in the satellite-diagnosed aerosol effect on the cloud liquid water path, Atmospheric Chemistry and Physics, 25, 1533–1543, https://doi.org/10.5194/acp-25-1533-2025, 2025.

Kurihana, T., Foster, I., Willett, R., Jenkins, S., Kurihana, K., Werman, R., Barros Lourenco, R., Neo, C., and Moyer, E.: Cloud classification with unsupervised deep learning, in: Proceedings of the 9th International Workshop on Climate Informatics: CI 2019, 2019.

Kurihana, T., Moyer, E. J., and Foster, I. T.: AICCA: AI-Driven Cloud Classification Atlas, Remote Sensing, 14, 5690, 2022.

Loveridge, J. and Di Girolamo, L.: Do Subsampling Strategies Reduce the Confounding Effect of Errors in Bispectral Retrievals on Estimates of Aerosol Cloud Interactions?, Journal of Geophysical Research: Atmospheres, 129, e2023JD040 189, 2024.

Mather, J. H. and Voyles, J. W.: The Arm Climate Research Facility: A Review of Structure and Capabilities, Bulletin of the American Meteorological Society, 94, 377 – 392, https://doi.org/10.1175/BAMS-D-11-00218.1, 2013.

Merk, D., Deneke, H., Pospichal, B., and Seifert, P.: Investigation of the adiabatic assumption for estimating cloud micro-and macrophysical properties from satellite and ground observations, Atmospheric Chemistry and Physics, 16, 933–952, 2016.

MODIS Characterization Support Team: MODIS 500m Calbrated Radiance Product, https://doi.org/10.5067/MODIS/MOD02HKM.061, 2017.

Mülmenstädt, J., Gryspeerdt, E., Dipu, S., Quaas, J., Ackerman, A. S., Fridlind, A. M., Tornow, F., Bauer, S. E., Gettelman, A., Ming, Y., Zheng, Y., Ma, P.-L., Wang, H., Zhang, K., Christensen, M. W., Varble, A. C., Leung, L. R., Liu, X., Neubauer, D., Partridge, D. G., Stier, P., and Takemura, T.: General circulation models simulate negative liquid water path–droplet number correlations, but anthropogenic aerosols still increase simulated liquid water path, Atmospheric Chemistry and Physics, 24, 7331–7345, https://doi.org/10.5194/acp-24-7331-2024, 2024.

Painemal, D., Smith Jr, W. L., Gupta, S., Moore, R., Cairns, B., McFarquhar, G. M., and O'Brien, J.: Can we rely on satellite visible/infrared microphysical retrievals of boundary layer clouds in partially cloudy scenes? implications for climate research, Geophysical Research Letters, 52, e2024GL113 825, 2025.

Platnick, S., Ackerman, S., King, M., Meyer, K., Menzel, W., Holz, R., Baum, B., and Yang, P.: MODIS atmosphere L2 cloud product (06_L2), NASA MODIS Adaptive Processing System, Goddard Space Flight Center, https://doi.org/http://dx.doi.org/10.5067/MODIS/MOD06_L2.061, 2015.

Platnick, S., Meyer, K. G., King, M. D., Wind, G., Amarasinghe, N., Marchant, B., Arnold, G. T., Zhang, Z., Hubanks, P. A., Holz, R. E., Yang, P., Ridgway, W. L., and Riedi, J.: The MODIS Cloud Optical and Microphysical Products: Collection 6 Updates and Examples From Terra and Aqua, IEEE Transactions on Geoscience and Remote Sensing, 55, 502–525, https://doi.org/10.1109/TGRS.2016.2610522, 2017.

Pruppacher, H. R. and Klett, J. D.: Microphysics of Clouds and Precipitation., 1997.

Quaas, J., Ming, Y., Menon, S., Takemura, T., Wang, M., Penner, J. E., Gettelman, A., Lohmann, U., Bellouin, N., Boucher, O., et al.: Aerosol indirect effects–general circulation model intercomparison and evaluation with satellite data, Atmospheric Chemistry and Physics, 9, 8697–8717, 2009.

Rasp, S., Schulz, H., Bony, S., and Stevens, B.: Combining Crowdsourcing and Deep Learning to Explore the Mesoscale Organization of Shallow Convection, Bulletin of the American Meteorological Society, 101, E1980 – E1995, https://doi.org/10.1175/BAMS-D-19-0324.1, 2020.

Riihimaki, L., McFarlane, S., and Sivaraman, C.: Droplet number concentration value-added product, Tech. rep., DOE ARM Climate Research Facility, Washington, DC (United States), 2021.

Segal Rozenhaimer, M., Nukrai, D., Che, H., Wood, R., and Zhang, Z.: Cloud Mesoscale Cellular Classification and Diurnal Cycle Using a Convolutional Neural Network (CNN), Remote Sensing, 15, https://doi.org/10.3390/rs15061607, 2023.

Stevens, B., Bony, S., Brogniez, H., Hentgen, L., Hohenegger, C., Kiemle, C., L'Ecuyer, T. S., Naumann, A. K., Schulz, H., Siebesma, P. A., et al.: Sugar, gravel, fish and flowers: Mesoscale cloud patterns in the trade winds, Quarterly Journal of the Royal Meteorological Society, 146, 141–152, 2020.

Tang, S., Varble, A. C., Fast, J. D., Zhang, K., Wu, P., Dong, X., Mei, F., Pekour, M., Hardin, J. C., and Ma, P.-L.: Earth System Model Aerosol–Cloud Diagnostics (ESMAC Diags) package, version 2: assessing aerosols, clouds, and aerosol–cloud interactions via field campaign and long-term observations, Geoscientific Model Development, 16, 6355–6376, https://doi.org/10.5194/gmd-16-6355-2023, 2023.

Tian, J., Comstock, J., Geiss, A., Wu, P., Silber, I., Zhang, D., Kooloth, P., and Feng, Y.-C.: Mesoscale cellular convection detection and classification using convolutional neural networks: Insights from long-term observations at ARM Eastern North Atlantic site, Journal of Geophysical Research: Machine Learning and Computation, 2, e2024JH000 486, 2025.

Twomey, S.: The Influence of Pollution on the Shortwave Albedo of Clouds, Journal of Atmospheric Sciences, 34, 1149 – 1152, https://doi.org/10.1175/1520-0469(1977)034<1149:TIOPOT>2.0.CO;2, 1977.

Varble, A. C., Ma, P.-L., Christensen, M. W., Mülmenstädt, J., Tang, S., and Fast, J.: Evaluation of liquid cloud albedo susceptibility in E3SM using coupled eastern North Atlantic surface and satellite retrievals, Atmospheric Chemistry and Physics, 23, 13 523–13 553, https://doi.org/10.5194/acp-23-13523-2023, 2023.

Wang, S., Wang, Q., and Feingold, G.: Turbulence, condensation, and liquid water transport in numerically simulated nonprecipitating stratocumulus clouds, Journal of the atmospheric sciences, 60, 262–278, 2003.

Wood, R.: Stratocumulus Clouds, Monthly Weather Review, 140, 2373 – 2423, https://doi.org/https://doi.org/10.1175/MWR-D-11-00121.1, 2012.

Wood, R. and Hartmann, D. L.: Spatial Variability of Liquid Water Path in Marine Low Cloud: The Importance of Mesoscale Cellular Convection, Journal of Climate, 19, 1748 – 1764, https://doi.org/10.1175/JCLI3702.1, 2006.

https://doi.org/10.5194/egusphere-2026-4426
Preprint. Discussion started: 31 July 2026

Wu, P., Dong, X., Xi, B., Tian, J., and Ward, D. M.: Profiles of MBL cloud and drizzle microphysical properties retrieved from ground-based observations and validated by aircraft in situ measurements over the Azores, Journal of Geophysical Research: Atmospheres, 125, e2019JD032 205, 2020.

Xue, H. and Feingold, G.: Large-Eddy Simulations of Trade Wind Cumuli: Investigation of Aerosol Indirect Effects, Journal of the Atmospheric Sciences, 63, 1605 – 1622, https://doi.org/10.1175/JAS3706.1, 2006.

Yuan, T.: Understanding Low Cloud Mesoscale Morphology with an Information Maximizing Generative Adversarial Network, Earth arXiv, https://doi.org/10.31223/osf.io/gvebt, 2019.

Yuan, T., Song, H., Wood, R., Mohrmann, J., Meyer, K., Oreopoulos, L., and Platnick, S.: Applying deep learning to NASA MODIS data to create a community record of marine low-cloud mesoscale morphology, Atmospheric Measurement Techniques, 13, 6989–6997, 2020.

Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S.: Barlow Twins: Self-Supervised Learning via Redundancy Reduction, in: Proceedings of the 38th International Conference on Machine Learning, edited by Meila, M. and Zhang, T., vol. 139 of *Proceedings of Machine Learning Research*, pp. 12 310–12 320, PMLR, https://proceedings.mlr.press/v139/zbontar21a.html, 2021.

Zelinka, M. D., Klein, S. A., and Hartmann, D. L.: Computing and Partitioning Cloud Feedbacks Using Cloud Property Histograms. Part I: Cloud Radiative Kernels, Journal of Climate, 25, 3715 – 3735, https://doi.org/10.1175/JCLI-D-11-00248.1, 2012.

Zhang, D., Riihimaki, L. D., Gaustad, K. L., and Turner, D. D.: MWRRETV2 Value-Added Product Report: The Retrieval of Liquid Water Path and Precipitable Water Vapor from Microwave Radiometer–3-Channel (MWR3C) Data Sets, Tech. rep., DOE Office of Science Atmospheric Radiation Measurement (ARM) Program . . . , 2020.

Zhang, Z. and Platnick, S.: An assessment of differences between cloud effective particle radius retrievals for marine water clouds from three MODIS spectral bands, Journal of Geophysical Research: Atmospheres, 116, 2011.