

MT-HCCAR: Multi-Task Deep Learning with Hierarchical Classification and Attention-based Regression for Cloud Property Retrieval

Xingyan Li^{1[0009-0001-2598-2296]}, Andrew M. Sayer^{2,3[0000-0001-9149-1789]}, Ian T. Carroll^{2,3}, Xin Huang⁴, and Jianwu Wang^{1[0000-0002-9933-1170]} (✉)

¹ Department of Information Systems, University of Maryland, Baltimore County, Baltimore, MD, USA {xingyanli,jianwuw}@umbc.edu

² NASA Goddard Space Flight Center, Greenbelt, MD, USA

³ Goddard Earth Sciences Technology and Research (GESTAR) II, UMBC, Baltimore, MD, USA

⁴ Dept of Computer & Information Sciences, Towson University, Towson, MD, USA

Abstract. In Earth science, accurate retrieval of cloud properties including cloud masking, cloud phase classification, and cloud optical thickness (COT) prediction is essential in atmospheric and environmental studies. Conventional methods rely on distinct models for each sensor due to their unique spectral characteristics. Recently, machine/deep learning has been embraced to extract features from satellite datasets, yet existing approaches lack architectures capturing hierarchical relationships among tasks. Additionally, given the spectral diversity among sensors, developing models with robust generalization capabilities remains challenging for related research. There is also a notable absence of methods evaluated across different satellite sensors. In response, we propose MT-HCCAR, an end-to-end deep learning model employing multi-task learning. MT-HCCAR simultaneously handles cloud masking, cloud phase retrieval (classification tasks), and COT prediction (a regression task). It integrates a hierarchical classification network (HC) and a classification-assisted attention-based regression network (CAR), enhancing precision and robustness in cloud labeling and COT prediction. Experimental evaluations, including comparisons with baseline methods and ablation studies, demonstrate that MT-HCCAR achieves optimal performance across various evaluation metrics and satellite datasets.

Keywords: Cloud property retrieval · Remote sensing · Deep learning.

1 Introduction

Clouds are integral components of the Earth system, wielding substantial influence over the planet’s energy dynamics, climate regulation, and the hydrological cycle [24]. Satellites have long been an indispensable tool to help us understand our environment. A prominent category of these instruments, commonly referred to as imagers, passively collect measurements of the Earth across various

combinations of ultraviolet (UV), visible (VIS), near-infrared (NIR), shortwave infrared (SWIR), and thermal infrared (TIR) wavelength ranges. These measurements of reflected solar and/or emitted thermal radiation in different spectral bands undergo routine processing by algorithms to convert them into geophysical parameters of interest (atmospheric and surface characteristics) in a process known as “retrieval”. This is typically done pixel-by-pixel for each satellite image. For clouds, several key properties are targeted, including a cloud mask (which distinguishes cloud-covered pixels from cloud-free ones), thermodynamic phase (indicating whether the cloud comprises liquid water droplets or ice crystals), and cloud optical thickness (COT), which is a measure of both the amount of light scattered by a cloud and the quantity of liquid or ice within it. The routine retrieval of geophysical parameters is essential for advancing our understanding of Earth’s climate [13].

One practical motivation for our work is NASA’s Plankton, Aerosol, Cloud, ocean Ecosystem (PACE) mission [28], which launched in February 2024. Existing NASA cloud masking algorithms are not directly applicable to PACE’s main sensor, called OCI [5], due to different spectral bands compared to algorithms developed for existing sensors. OCI has similarities with some previous space-borne imagers such as MODIS [3], VIIRS [6], and ABI [1] sensor types. However, two key differences are that 1) OCI has continuous hyperspectral coverage in the UV-NIR, plus discrete SWIR bands, while the others are only multi-spectral (up to a dozen or so discrete relevant bands); and 2) OCI lacks TIR bands. Some of the most commonly used cloud masking tests for those sensors are therefore not applicable, and adapting a subset of those tests would miss out on OCI’s unique abilities, so a new approach is warranted.

Numerous studies have explored machine learning techniques to extract cloud properties from satellite sensor data. These include both retrieval of a single cloud property [17, 25, 12], and simultaneous retrievals of multiple cloud properties [26, 14]. However, challenges persist in this area. To begin with, it’s unclear how incorporating atmospheric domain knowledge, such as physical relationships between cloud properties, alongside advanced machine learning techniques, enhances retrieval accuracy. Furthermore, despite the deployment of numerous satellite sensors for similar cloud retrieval tasks (as discussed in Section 3), the generalizability, especially from an Earth science perspective, of employing a unified machine learning architecture across different sensors remains uncertain.

To address the above challenges we introduce MT-HCCAR, an end-to-end Multi-Task Learning (MTL) model for cloud masking, cloud phase prediction, and COT regression. Our contributions are fourfold. First, we incorporate hierarchical classification to capture the hierarchical relationship between cloud masking and cloud phase classification, enhancing prediction performance as demonstrated in the comparison experiment with baseline methods and ablation study. Second, we employ a cross-attention module to improve COT regression accuracy by leveraging similarities across the classification and regression networks. Third, we conduct quantitative experiments to demonstrate the advantages of our MT-HCCAR model over state-of-the-art baselines and its generalizability

across three different satellite sensors. Finally, we perform quantitative and qualitative evaluation of performance variations across sensors from an Earth science perspective. Our implementation is publicly available at GitHub [2].

2 Related Work

In recent years, various machine learning techniques, such as Random Forest (RF), Multi-Layer Perceptron (MLP), and Convolutional Neural Network (CNN), have been employed for retrieving different cloud properties. Among them, most of the work targets cloud detection [17, 25], cloud phase [11], and cloud thickness [21]. While these machine learning approaches for cloud property retrieval predominantly leverage spectral features, the main limitations lie in two aspects: 1) many studies do not consider background knowledge, such as the task order (e.g., cloud mask prediction, cloud phase prediction, and COT prediction), and 2) several studies conduct different tasks independently, lacking knowledge sharing between classification and regression tasks. These limitations underscore the need for more integrated and informed methodologies in cloud property retrieval studies.

Multi-task learning (MTL) is proving valuable in Earth science and remote sensing by jointly enhancing performance across diverse remote-sensing tasks through shared features. The studies have shown that classification and regression tasks can be implemented together to improve model performance [15, 7]. There are also studies conducting ML learning specifically for cloud property retrieval. Yang et al. [29] developed an MLP-based method to retrieve cloud macrophysical parameters (cloud mask, cloud top temperature, and cloud top height) using Himawari-8 satellite data. Wang et al. [26] proposed TIR-CNN based on the U-Net model to retrieve cloud properties including cloud mask, COT, effective particle radius (CER), and cloud top height (CTH) from thermal infrared radiometry. The architecture consists of encoding and decoding layers, convolutional blocks, batch normalization layers, and leaky Rectified Linear Units (ReLU). The results of applying the model to thermal infrared spectral data from MODIS are used to compare model performance for daytime and nighttime data.

While these above studies offer valuable insights that are beneficial to our design, our work cannot be compared with most of them directly via experiments. This is because the OCI data we use does not have the same spectral bands, such as TIR bands, with those methods used. Consequently, we will not compare our model with architectures employing convolutional layers, such as CNN.

3 Problem Statement and Data Simulation

Cloud property determination requires two operations: simulation and retrieval. Simulation is a forward process mapping from the Earth’s state (cloud and surface properties) to satellite direct observations (the reflected light signal the

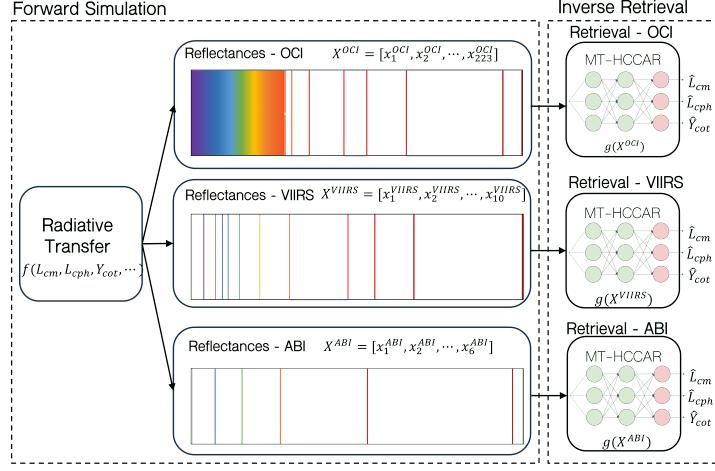


Fig. 1. Illustration of forward simulation and retrieval. \mathbf{x}_M : the spectral and physical features. $f(\cdot)$, $g(\cdot)$: the RT simulation algorithm and the retrieval algorithm. \mathbf{L}_{cm} , \mathbf{L}_{cph} , \mathbf{Y}_{COT} : the labels of cloud mask, cloud phase, and COT values, respectively.

satellites see). Retrieval is an inverse process, determining the geophysical quantities from the satellite observations. Figure 1 illustrates these two processes with colors showing differences in the wavelength band of the three sensors.

3.1 Radiative Transfer Simulation

The simulation process provides training data for our cloud property retrieval model. This study was conducted before the launch of PACE, and the spacecraft has not entered routine operations. Therefore, simulated datasets for OCI, VIIRS, and ABI were utilized to evaluate the sensors and architectures against a standardized source, enabling direct comparisons. This approach ensures consistency as the satellites are tested on identical simulated scenes with the same viewing geometries, which would not be feasible using real data.

Radiative transfer (RT) models describe the scattering and absorption processes that affect the propagation of light through Earth's atmosphere and surface. RT is the forward model to map from surface and atmospheric properties to spectral intensities seen by the space-borne sensors. We account for realistic variations in the properties of clouds (phase, microphysical, optical, and vertical structure), the atmospheric profile (aerosol, gas, Rayleigh scattering, temperature, and pressure), surface reflectance, and solar/observation geometry. Further details are as in Sayer et al. [20], with two modifications: 1) simulating cloud-free conditions as well as single-layer clouds, and 2) utilizing 20 different surface reflectance classes included in the libRadtran RT model [10] (including various land surface spectra, water, and snow/sea ice). The simulation does not include multi-layer or mixed-phase cloud systems, aligning with most satellite retrieval

processing algorithms. However, future work may extend the network to include such systems. We generate a dataset of 250,000 simulations for model training and evaluation. We simulate from the UV to SWIR and convolve these simulations with the solar spectrum from [9] and sensor relative spectral response functions for OCI [5], ABI [1], and VIIRS [6] in order to generate the spectral top of atmosphere reflectance signal that the instruments would observe. This provides our simulated observations and reference truth (cloud classification and COT), along with band centers for each sensor. For **OCI**, there are 233 bands in total, including 226 hyperspectral bands ranging from 350 to 890 nm with 2.5 nm spacing, and seven discrete NIR/SWIR bands centered near 940, 1040, 1250, 1378, 1620, 2130, and 2260 nm. **VIIRS** has 10 bands centered near 412, 445, 488, 555, 672, 865, 1240, 1380, 1610, and 2250 nm. **ABI** has 6 bands centered near 471, 640, 860, 1370, 1600, and 2200 nm, respectively.

Each ABI band has a close analog in VIIRS, and similarly, each VIIRS band has a counterpart in OCI, thus forming a trio representing increasing complexity. Another commonly-used sensor, MODIS, has similar spectral coverage to VIIRS so we do not include it. Both ABI and VIIRS have TIR bands, but these are omitted from consideration as our primary focus is on developing a new methodology for OCI, which lacks TIR bands. Application to ABI and VIIRS demonstrates the broader applicability of our new architectures.

3.2 Cloud Property Retrieval

Our primary objective is to accurately model the base-10 logarithm of COT (throughout the remaining text, COT means this logarithm unless explicitly described as the “original” COT) for pixels labeled as cloudy. We work in log space for COT as this has a more linear relationship with the brightness seen by satellite than the original COT. Simultaneously, the model should be able to accurately classify the cloud phase [cloudy, cloud-free, cloudy-liquid, cloudy-ice] for each pixel in the dataset to aid COT prediction. Therefore, our problem consists of two tasks: 1) a classification task to predict cloud mask and phase, and 2) a regression task to predict COT values. Assuming that the labels are \bar{C} = ‘cloud-free’, C = ‘cloudy’, CL = ‘cloudy-liquid’, CI = ‘cloudy-ice’, and bold notations indicate arrays, the details of our study are outlined below.

Model input. 1) **Input features** are represented as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$, where M is the dimension of the available features. The exact input feature variables include: i) surface pressure in millibar (mbar), ii) total column water vapor content in millimeters (mm), iii) total column ozone content in Dobson units (DU), iv) types of Earth surface with 4 categories (land, snow, desert, and ocean/water), v) top of atmosphere reflectance at different wavelengths collected by each spaceborne sensor, vi) viewing zenith angle, solar zenith angle, and the relative azimuth angle, in degree. 2) **Cloud mask/phase labels** \mathbf{l}^{cls} are used to train the classification task of the model. The set of possible label values is $l = \{C, \bar{C}, CL, CI\}$. There is a hierarchical relationship between the labels, as pixels covered by liquid cloud or ice cloud are both cloudy pixels. Also, there is no coexistence between liquid cloud and ice cloud in our data. Thus, CL and

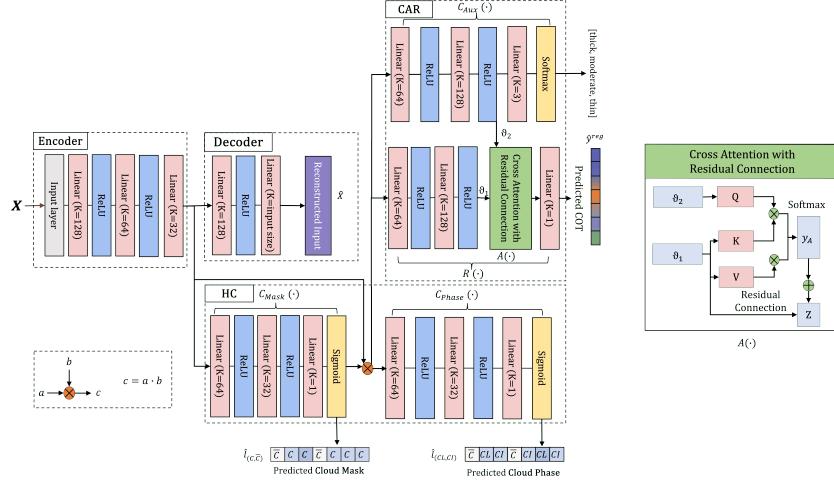


Fig. 2. MT-HCCAR: an end-to-end MTL model with hierarchical classification (HC) and cross attention assisted regression (CAR). The HC sub-network comprises the cloud masking module C_{Mask} and the cloud phase classification module C_{Phase} . The CAR sub-network consists of the auxiliary coarse classification module C_{Aux} , the cross attention module A , and the regression module R . On the right is the structure of A .

CI can be two subclasses of label C = ‘cloudy’. 3) **True COT value y^{cot}** are used to train the regression task of the model. COT values are not available for pixels with no cloud cover. That is, if $l_i^{cls} = 0$, then $y_i^{cot} = N/A$.

Model output. 1) Predicted cloud mask/phase class \hat{l}^{cls} with probabilities \mathbf{u} of each pixel belonging to each of the four classes. 2) Predicted COT value \hat{y}^{cot} for each pixel. 3) Model architecture M with parameters β . The predicted values are generated by the model, which is $[\hat{y}^{cot}, \hat{l}^{cls}] = M(\mathbf{X}|\beta)$.

4 MT-HCCAR Model

Our objective is to train a deep learning model to accomplish two tasks: 1) the classification of cloud mask and phase for each pixel based on its reflectance values, and 2) the subsequent prediction of COT values for pixels classified as cloudy. Toward this objective, our proposed MT-HCCAR model is illustrated in Figure 2 and the subsections below will each of the modules.

4.1 Encoder-Decoder Sub-Network

An encoder-decoder sub-network, containing an encoder module and a decoder module, is employed in our model to learn latent feature parameters that can be shared for other learning tasks. To integrate the COT regression task and cloud mask/phase prediction task, we adopt the soft-parameter sharing approach of

MTL where the shared parameters are derived by encoder-decoder. Previous studies [22, 16] have used similar encoder-decoder techniques to transform original features into more relevant features for better downstream prediction accuracy. In our model, the decoder reconstructs the input feature \mathbf{X} into $\hat{\mathbf{X}}$, and a loss is between \mathbf{X} and $\hat{\mathbf{X}}$ is minimized during training to minimize the distortion of input feature throughout feature extraction layers.

The shared encoder is formed by the first three dense layers, each of which is followed by a ReLU activation function. This feature extractor is a wide-to-narrow structure where the number of kernels at each layer decreases gradually. The decoder and reconstruction branch is narrow-to-wide and further improves the performance of the encoder to learn features for different tasks.

4.2 Hierarchical Classification (HC) Sub-Network

Utilizing the shared parameters acquired by the encoder-decoder sub-network, MT-HCCAR incorporates a hierarchical classification sub-network. This sub-network comprises a cloud mask classification module and a cloud phase classification module, facilitating cloud mask and phase predictions. To enhance the physical interpretability of the model, the architecture is crafted to mirror human cognitive processes in understanding cloud labels. Fundamentally, to predict cloud phase and COT values from an Earth science perspective, the model must discern between cloudy and non-cloudy pixels before further classifying liquid-phase and ice-phase pixels and predicting COT values for cloudy pixels.

Hierarchical classification has the capability to categorize instances according to label levels, forming a tree-like structure where each label functions as a node on the tree [18]. The HC network used in MT-HCCAR consists of two classifiers C_{Mask} and C_{Phase} , the first of which distinguishes cloudy pixels from the cloud-free pixels, and the output is uncertainties for the two labels $[\hat{\mathbf{u}}^C, \hat{\mathbf{u}}^{\bar{C}}]$. The output uncertainties are binarized to $[\hat{\mathbf{I}}^C, \hat{\mathbf{I}}^{\bar{C}}]$ by a threshold of 0.5. Before the beginning of the second classifier C_{Phase} input feature \mathbf{X} is multiplied by the predicted cloud label, $\hat{\mathbf{y}}^C$, to filter out cloud-free pixels. Then C_{Phase} classifies liquid-cloud and ice-cloud and the sigmoid uncertainties $[\hat{\mathbf{u}}^{CL}, \hat{\mathbf{u}}^{CI}]$ are transformed in the same way to binary cloud-phase labels $[\hat{\mathbf{I}}^{CL}, \hat{\mathbf{I}}^{CI}]$.

4.3 Classification Assisted Regression Sub-Network based on Cross Attention Mechanism (CAR)

Taking inputs from the encoded feature set and cloud mask classification results as input, our regression sub-network CAR is used for further downstream regression prediction of a cloudy pixel’s COT value. Two novel efforts were taken to improve COT prediction accuracy. First, instead of a direct regression module, an auxiliary coarse classification module is added to predict which sub-range of COT each pixel falls into. Second, a residual-based cross-attention mechanism inspired by [23] is developed to enable the regression module and the auxiliary classification module to share relevant correlations and insights.

As illustrated in Figure 2, there is a connection between internal features from the auxiliary classifier and those from the regression network, facilitated by a residual-based cross-attention module. The close alignment in tasks between the auxiliary classifier and the regression network enables the cross-attention mechanism to selectively attend to features relevant to one task while simultaneously executing the other. This contrasts with relationships involving the regression task and other tasks, such as cloud phase classification, where associations are forged through joint optimization of their respective losses during training iterations, rather than through the shared utilization of internal features.

Our auxiliary coarse classification involves discretizing continuous values into coarse groups, serving as a preprocessing step before regression to align pixels with similar characteristics in the feature space. Moreover, the auxiliary classifier assigns greater importance to the regression task during model training to get accurate COT predictions. In MT-HCCAR, an auxiliary classifier is employed to categorize continuous COT values into three distinct levels: thin cloud, moderate cloud, and thick cloud. Specifically, we define thin cloud pixels as those with logarithmic COT values within the range of $[-1.5, 0]$, moderate cloud in $[0, 1]$, and thick cloud in $[1.0, 2.5]$. Given the analogous nature of the tasks in COT coarse classification and COT regression, the features extracted by the auxiliary classifier are integrated with those obtained from the regression network.

To learn the joint features more effectively, we utilize cross-attention layers to enhance the integration of features from both regression layers and coarse classification layers, facilitating deeper feature learning and more cohesive feature amalgamation. Based on the cross-attention mechanism, adding a residual connection guarantees that the module has a stable output and related research such as Wang et al. [27] propose a non-local operation network using a residual connection to insert blocks to the network. By using the residual connection in our module A , the output is $A(\vartheta_1, \vartheta_2) = W_z y_A(\vartheta_1, \vartheta_2) + \vartheta_1$, where as shown in Figure 2 ϑ_1 is the feature from the regression network, ϑ_2 is the feature from the auxiliary classification module, W_z is the weight matrix to map y_A to ϑ_2 , and y_A denotes the computation within cross-attention mechanism. The calculation of y_A involves three variables: a query (Q), a key (K) and a value (V) [23]. They are calculated based on corresponding weights matrices W_Q , W_K , W_V learned through model training. $Q = W_Q \vartheta_2$, $K = W_K \vartheta_1$, and $V = W_V \vartheta_1$. Then the y_A is $y_A = \text{Softmax}(QK^T)V$.

4.4 Model Training of MT-HCCAR

The loss function $\mathcal{L}_{MT-HCCAR}$ for training MT-HCCAR model is formulated as the weighted sum of four components: a hierarchical classification loss \mathcal{L}_{HC} , a regression loss \mathcal{L}_{CAR} , a reconstruction loss \mathcal{L}_{Rec} , and a Lasso regularization loss \mathcal{L}_{Lasso} . That is: $\mathcal{L} = \mathcal{L}_{HC} + \mathcal{L}_{CAR} + \mathcal{L}_{Rec} + \mathcal{L}_{Lasso}$. The four loss components in the loss function are calculated using different rules.

Among the four components, The reconstruction loss \mathcal{L}_{Rec} and the Lasso regularization loss component \mathcal{L}_{Lasso} can be directly calculated. The reconstruction

loss \mathcal{L}_{Rec} describes the difference between the input features \mathbf{X} and the reconstruction of the input generated through the encoder and the decoder. The loss function is a mean square error (MSE) between $\mathbf{D}(\mathbf{X}) = \hat{\mathbf{X}}$ and \mathbf{X} , which means $\mathcal{L}_{Rec} = \sum_{i=1}^n (\mathbf{X}_i - \hat{\mathbf{X}}_i)^2$. The Lasso regularization loss \mathcal{L}_{Lasso} is an additional penalty to regularize the training process by increasing the sparsity of the model with the Lasso regularization loss. $\mathcal{L}_{Lasso} = \lambda \sum_{p=1}^P |\beta_p|$.

The other two components, HC loss and CAR loss, are both weighted sums of losses from the supplement modules. The HC loss is the summation of Binary Cross Entropy loss from cloud masking module C_{Mask} and cloud phase classification module C_{Phase} . The HC loss function is $\mathcal{L}_{HC} = \mathcal{L}_{C_{Mask}} + \mathcal{L}_{C_{Phase}}$. We have a sigmoid output of cloudy and non-cloud uncertainties $[\hat{u}^C, \hat{u}^{\bar{C}}]$ from cloud mask classifier C_{Mask} and sigmoid output of liquid cloud and ice cloud uncertainties $[\hat{u}^{CL}, \hat{u}^{CI}]$ from cloud phase classifier C_{Phase} , then for the HC loss, $\mathcal{L}_{C_{Mask}}$ and $\mathcal{L}_{C_{Phase}}$ are calculated as $\mathcal{L}_{C_{Mask}} = -\frac{1}{N} \sum_{i=1}^N l_i^C \cdot \log(\hat{u}_i^C) + l_i^{\bar{C}} \cdot \log(\hat{u}_i^{\bar{C}})$ and $\mathcal{L}_{C_{Phase}} = -\frac{1}{N} \sum_{i=1}^N \hat{u}_i^C l_i^{CL} \cdot \log(\hat{u}_i^C \hat{u}_i^{CL}) + \hat{u}_i^C l_i^{CI} \cdot \log(\hat{u}_i^C \hat{u}_i^{CI})$. The CAR loss \mathcal{L}_{CAR} is the summation of Cross Entropy loss from the auxiliary classifier C_{Aux} and $l1$ loss from regression. That is, $\mathcal{L}_{CAR} = \mathcal{L}_R + \mathcal{L}_{C_{Aux}}$. We have predicted COT values and \hat{y}^{cot} and sigmoid output of thickness group uncertainties $[\hat{u}^{thin}, \hat{u}^{mod}, \hat{u}^{thick}]$ from the auxiliary classifier, then the two components are $\mathcal{L}_R = \sum_{i=1}^N |y_i^{cot} - \hat{y}_i^{cot}|$ and $\mathcal{L}_{C_{Aux}} = -\sum_{i=1}^N \sum_{c \in [thin, mod, thick]} l_i^c \ln \hat{u}_i^c$.

5 Experiments

To evaluate our model’s performance, we compare MT-HCCAR with baseline methods and conduct an ablation study across OCI, VIIRS, and ABI datasets.

Given that our dataset comprises independent pixels, we choose two baseline methods from prior research with comparable task objectives and data formats for our experiments: 1) **Chen et al. (2018)** [8]: An MLP based method for cloud property retrieval. We use 1 hidden layer with 10 nodes, which are the same as the authors used. 2) **Liu et al. (2022)** [17]: An RF based method for cloud detection. We apply one RF to cloud masking and the same RF to cloud phase classification, respectively. The parameters involved in the method are $ntrees = 100$ and $mdepth = 10$.

Besides the above two baseline methods, four ablation study models are also used to evaluate the effectiveness of different modules in our model. The comparison of the four models and our MT-HCCAR is shown in Table 1. The detailed differences can be found in the supplementary material.

5.1 Experiment Setup

Our model was implemented using the Python deep-learning library PyTorch. All baseline models and proposed models were tested on a single GPU on Kaggle. The dataset comprises satellite data simulation of instruments OCI, VIIRS, and ABI. The three datasets encompassing $N = 250,000$ instances.

Table 1. Comparison of models in the ablation study. STL: single-task learning.

MODEL	STRUCTURE	CLASSIFICATION	REGRESSION
SEQ	STL	Regular	Regular
MT-CR	MTL	Regular	Regular
MT-HCR	MTL	HC	Regular
MT-HCCR	MTL	HC	CAR without Attn.
MT-HCCAR	MTL	HC	CAR

We split data into 62.5%, 22.5%, and 10.0% for training, validation, and test sets, respectively. All models are trained with the same hyperparameters including learning rate = $1e^{-5}$, batch size = 64, and training epochs = 500. All experiment results are mean values of metrics from a 10-fold cross-validation.

Simulations of OCI, VIIRS, and ABI are utilized instead of observations as our labeled training dataset principally because there is no suitable comprehensive reference truth dataset. While the standard retrieval products from these sensors could be used to train networks, for example, these products have known limitations so there would be a risk of the model training to the artifacts in these products. In a similar vein, comparing the results of applying the trained NN to real satellite observations of these standard products will be instructive to get a general sense of reasonableness but not to gauge their absolute performance. It is expected that, by using a more realistic training set of simulations than was used to develop the physically-based retrievals (which include many simplifying assumptions out of the computational necessities from decades ago when these approaches were developed), the NN model should outperform them.

5.2 Evaluation Metrics

We use three types of metrics to evaluate our work, including metrics for the classification task, metrics for the regression task, and metrics from an Earth science perspective.

Metrics for classification performance. We use Accuracy and Average precision to evaluate classification performance. Cloud masking accuracy (ACC_{bi}) is the fraction of correct predictions of two big categories [cloudy, cloud-free] by our model. The area under the precision-recall curve ($AU(\overline{PRC})_c$) is calculated for each label. The weighted area under the precision-recall curve ($AU(\overline{PRC})_w$) is the weighted mean of precisions across all labels at each threshold h for class i , with the weight as subtracted between recall at threshold h and the recall at threshold $h - 1$, where the number of thresholds is close to infinity. The weighted precision and recall are: $R_w^h = \frac{\sum_l TP_c^h}{\sum_l TP_c^h + \sum_l FN_c^h}$ and $P_w^h = \frac{\sum_l TP_c^h}{\sum_l TP_c^h + \sum_l FP_c^h}$, then $AU(\overline{PRC})_w$ is: $AU(\overline{PRC})_w = \sum_h (R_w^h - R_w^{h-1})P_w^h$.

Metrics for regression performance. We use MSE and R^2 to evaluate regression performance. Mean squared error (MSE) is the average squared difference between the true and predicted COT values. $MSE = \sum_{i=0}^n (y_i^{cot} - \hat{y}_i^{cot})^2$.

The coefficient of determination (R^2) measures how close the predicted value is to the true value. $R^2 = 1 - \frac{\sum_{i=0}^n (y_i^{cot} - \hat{y}_i^{cot})^2}{\sum_{i=0}^n (y_i^{cot} - \bar{y}^{cot})^2}$, where \bar{y}^{cot} is the mean value of all true COT values $[y_1^{reg}, y_2^{reg}, \dots, y_n^{reg}]$.

Earth science metrics. Besides the above metrics, we also evaluate how a machine learning model could be used for actual satellite missions. The fraction of pixels meeting PACE goals (FMG) is an evaluation metric defined by the PACE Validation Plan [4] based on scientific requirements and expectations. For each pixel, the relative error $Error_i = |\frac{y^{(cot)_i} - \hat{y}^{(cot)_i}}{y^{(cot)_i}}|$. Then FMG represents the percentage of pixels whose relative error is less than 0.25 (for liquid clouds) or 0.35 (for ice clouds). The Validation Plan defines a predictive model with satisfactory performance as one where this goal is met for 65% of cloudy pixels. Note that this mission goal only applies to cases where the true original COT value is 5 or more ($\log_{10} \text{COT} > 0.7$).

5.3 Comparison with Baseline Models

Tables 2 and 3 present a comprehensive evaluation of our proposed model MT-HCCAR against two baseline methods, demonstrating superior performance across both classification and regression metrics. Specifically, in the assessment of classification tasks, ACC_{bi} pertains to cloud masking performance, while $AU(\overline{PRC})_w$ and $AU(\overline{PRC})_c$ quantify the performance of both cloud masking and cloud phase classification.

Upon close examination of Table 2, the RF classifier by [17] attains superior ACC_{bi} , $AU(\overline{PRC})_w$, and $AU(\overline{PRC})_c$ in comparison to the SEQ model and the simplest MTL based model, MT-CR, within the ablation study. However, MTL-based models featuring the HC module including MT-HCR, MT-HCCR, and MT-HCCAR, surpass the performance of the RF classifier in these metrics. Shifting the focus to the results of the COT regression task in Table 3, the MLP-based baseline method introduced by [8] produces results on par with the SEQ model but significantly inferior to MTL-based models. This finding emphasizes the effectiveness of MTL-based models for COT regression.

5.4 Ablation Study

The ablation study with five models from SEQ to MT-HCCAR for OCI, as depicted in the lower part of Tables 2 and 3, highlights the effectiveness of our model in improving both classification and regression performance. The comparative analysis demonstrates significant improvements across all metrics from SEQ to MT-CR, which shows the usefulness of the MTL structure. For MTL methods, MT-HCR outperforms MT-CR in binary classification (cloudy, cloud-free) and refines liquid and ice cloud phase classification, leading to enhancements in regression metrics. This proves the effectiveness of the HC module. Moreover, despite focusing on regression, the introduction of CAR in MT-HCCAR maintains the performance of classification tasks. The efficacy of the CAR module is further validated by comparing MT-HCCAR to MT-HCR, indicating improvements

Table 2. Model performance comparison with baseline models and ablation study - classification results. Models are trained on OCI.

MODEL	ACC _{bi}	AU(\overline{PRC}) _w	AU(\overline{PRC}) _c			
			Cloudy	Cloud-free	Liquid	Ice
[17]	0.968	0.985	0.955	0.949	0.951	0.965
[8]	0.942	0.638	0.974	0.882	0.385	0.348
SEQ	0.963	0.750	0.997	0.978	0.618	0.658
MT-CR	0.965	0.955	0.925	0.925	0.954	0.967
MT-HCR	0.975	0.993	0.999	0.989	0.982	0.992
MT-HCCR	0.982	0.995	0.999	0.992	0.985	0.993
MT-HCCAR	0.984	0.996	0.999	0.995	0.990	0.996

Table 3. Model performance comparison with baseline models and ablation study - regression results. Models are trained on OCI. The model [17] is not applicable for the regression task.

MODEL	MSE	MSE EACH LABEL	R ²		R ² EACH LABEL		FMG	
			Liquid	Ice	Liquid	Ice	Liquid	Ice
[8]	0.185	0.193	0.178	0.235	0.247	0.212	43.92%	60.64%
SEQ	0.126	0.133	0.118	0.440	0.274	0.548	4.00%	41.11%
MT-CR	0.062	0.075	0.049	0.728	0.661	0.788	44.46%	68.21%
MT-HCR	0.038	0.045	0.030	0.843	0.810	0.874	65.93%	82.28%
MT-HCCR	0.031	0.035	0.028	0.869	0.854	0.884	66.38%	82.43%
MT-HCCAR	0.026	0.032	0.021	0.891	0.869	0.914	67.96%	82.70%

in all metrics. Additionally, the adoption of the cross-attention module is substantiated by superior performance in MT-HCCR compared to MT-HCCAR, emphasizing its role in facilitating information exchange between hierarchical classification and regression.

To confirm the generalizability of our proposed model, we did a further ablation study with VIIRS and ABI. As shown in Table 4, the application of MTL-based models reveals a similar trend of performance enhancement when integrating HC and CAR modules. Notably, MT-HCCAR achieves the highest performance across both classification and regression tasks among the MTL models across OCI, VIIRS, and ABI datasets, underscoring the generalization capabilities of the introduced model MT-HCCAR. The detailed results of cross-validation and model selection are in the supplementary material.

To visually illustrate the improvements facilitated by the HC and CAR modules, Figure 3 presents a scatter plot showing the distribution of true COT values against predicted COT values for all instances in the test set. Integration of HC and CAR modules results in more accurate predictions compared to MT-CR and MT-HCR. Notably, the predicted Probability Density Function (PDF) aligns closely with the true distribution, indicating improved fidelity to actual

Table 4. Model performance for VIIRS and ABI datasets.

DATASET	MODEL	ACC _{bi}	AU(<i>PRC</i>) _w MSE	R ²	FMG (<i>CL</i> , <i>CI</i>)	
VIIRS	MT-CR	0.976	0.939	0.043	0.810	48.52% 72.73%
	MT-HCR	0.989	0.996	0.030	0.874	66.23% 82.62%
	MT-HCCAR	0.989	0.997	0.025	0.897	68.20% 82.69%
ABI	MT-CR	0.973	0.969	0.057	0.751	47.87% 71.81%
	MT-HCR	0.986	0.996	0.038	0.840	66.58% 82.55%
	MT-HCCAR	0.987	0.996	0.032	0.866	68.54% 83.34%

COT values. Additionally, scatter points cluster more tightly along the diagonal line, confirming the model’s enhanced precision when HC and CAR modules are incorporated into the model.

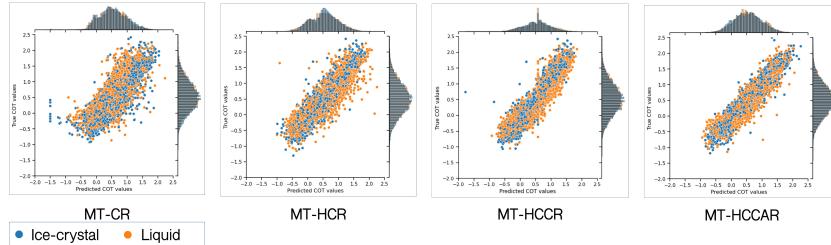


Fig. 3. Scatter plot with PDF distribution of true COT values (y-axis) and predicted COT values (x-axis) for the OCI dataset.

5.5 Earth Science Evaluation

Quality of performance. Satellite cloud masks are generally evaluated either through comparison against ground-based, airborne, or spaceborne observations from radar/lidar sensors that have much greater sensitivity for cloud detection than imagers in question. Accuracies tend to depend on the surface type and illumination conditions; for daytime scenes such as those simulated here, reported accuracies for NASA’s widely-used MODIS cloud mask are 0.850, 0.778, and 0.894. Chen et al. [8] developed an NN cloud mask for MODIS based on RT simulations and found an accuracy of 0.985 on training data but lower accuracies in the range of 0.739-0.885 on real data depending on region and time of year. Although real-world conditions are more complex than a simulation, the results suggest our models are competitive with current approaches. Note, that the cloud phase is less commonly validated in this way.

Obtaining a validation-grade measurement of COT presents significant challenges, with reference datasets lacking comprehensive coverage [19]. Consequently,

existing comparisons primarily focus on evaluating agreement among different remotely sensed datasets. Our network’s performance suggests its potential to fulfill the objectives outlined in the Validation Plan [4]. Notably, PACE’s allowance for larger uncertainty in ice cloud measurements stems from the anticipated difficulties arising from variations in ice crystal properties in real-world scenarios. However, our results show that the uncertainty in ice COT can be smaller than that for liquid COT. Traditional methods often assume specific ice crystal properties, resulting in increased uncertainty in retrieved ice COT. Our findings indicate that satellite measurements inherently contain information about these properties, which neural networks are proficient at learning.

Performance of different satellite sensors. We also compared how the models perform for different satellite sensors based on the results in Tables 2, 3 and 4. The classification tasks exhibit high accuracy across all sensors, with virtually indistinguishable performance from a scientific perspective. The fundamental nature of cloud masking in Earth science prompts many sensors to incorporate a common set of bands proven to be effective for this purpose, with additional bands often designed for diverse applications. For instance, OCI’s hyperspectral bands support the measurement of different ocean plankton species [28], revealing subtle spectral differences not readily discernible in multispectral data. In the regression task, OCI and VIIRS outperform ABI, indicating the utility of additional bands in predicting COT. The unexpectedly slightly superior performance of VIIRS over OCI may suggest that OCI’s substantially larger feature space makes finding an optimal solution during training more challenging or could be attributed to stochastic variation. Adding training epochs for OCI may lead to better results, considering the number of bands.

6 Conclusions

In this study, we present MT-HCCAR, an end-to-end multi-task learning model tailored for cloud property retrieval on a simulated OCI satellite dataset in the PACE project, tackling tasks including cloud masking, cloud phase classification, and COT prediction. The model is implemented on three sensors’ simulated datasets (OCI, VIIRS, and ABI), respectively, to examine its generalization. Comparative analyses against two baseline methods and ablation studies underscore the effectiveness of the HC module and the CAR module, enhancing performance in both classification and regression tasks. The ablation study establishes MT-HCCAR’s superior performance across different datasets and multiple evaluation metrics. The positive results affirm our model’s capability to address real-world challenges in cloud property retrieval and other multi-task applications. Future research endeavors will involve applying the model to spatial or temporal OCI images, co-located in space and time with VIIRS and ABI post-PACE launch, to assess its performance, and consistency, and enable detailed comparisons with deep learning models and non-machine-learning approaches.

Acknowledgments. This work is supported by a student fellowship from Goddard Earth Sciences Technology and Research (GESTAR) II, UMBC, grant OAC-1942714

from the National Science Foundation (NSF) and grant 80NSSC21M0027 from the National Aeronautics and Space Administration (NASA).

References

1. Advanced Baseline Imager (ABI) Sensor. <https://ncc.nesdis.noaa.gov/GOESR/ABI.php>, last accessed 2024-03-22
2. GitHub Repository of Multi-Task Deep Learning with Hierarchical Classification and Attention-based Regression for Cloud Property Retrieval (MT-HCCAR). <https://github.com/AI-4-atmosphere-remote-sensing/MT-HCCAR>, last accessed 2024-06-10
3. Moderate Resolution Imaging Spec- troradiometer (MODIS) Sensor. <https://modis.gsfc.nasa.gov/>, last accessed 2024-03-22
4. NASA PACE Validation Plan, 2020. https://pace.oceansciences.org/docs/PACE_Validation_Plan_14July2020.pdf, last accessed 2024-03-22
5. Ocean Color Instrument (OCI) Sensor. <https://pace.oceansciences.org/oci.htm>, last accessed 2024-03-22
6. Visible Infrared Imaging Radiometer Suite (VIIRS) Sensor. <https://ncc.nesdis.noaa.gov/VIIRS/>, last accessed 2024-03-22
7. Chen, J., Hu, W., Cao, D., Zhang, Z., Chen, Z.: A novel multi-task learning method with attention mechanism for wind turbine blades imbalance fault diagnosis. In: 2022 4th Asia Energy and Electrical Engineering Symposium (AEEES). pp. 857–862. IEEE (2022)
8. Chen, N., Li, W., Gatebe, C., Tanikawa, T., Hori, M., Shimada, R., Aoki, T., Stamnes, K.: New neural network cloud mask algorithm based on radiative transfer simulations. *Remote Sensing of Environment* **219**, 62–71 (12 2018)
9. Coddington, O.M., Richard, E.C., Harber, D., Pilewskie, P., Woods, T.N., Chance, K., Liu, X., Sun, K.: The tsis-1 hybrid solar reference spectrum. *Geophysical research letters* **48**(12), e2020GL091709 (2021)
10. Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., Pause, C., Dowling, T., et al.: The libradtran software package for radiative transfer calculations (version 2.0. 1). *Geoscientific Model Development* **9**(5), 1647–1672 (2016)
11. Guo, J., Yang, J., Yue, H., Tan, H., Hou, C., Li, K.: Cdnetv2: Cnn-based cloud detection for remote sensing imagery with cloud-snow coexistence. *IEEE Transactions on Geoscience and Remote Sensing* **59**(1), 700–713 (2020)
12. He, Q., Sun, X., Yan, Z., Fu, K.: Dabnet: Deformable contextual and boundary-weighted network for cloud detection in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–16 (2021)
13. Hollmann, R., Merchant, C.J., Saunders, R., Downy, C., Buchwitz, M., Cazenave, A., Chuvieco, E., Defourny, P., de Leeuw, G., Forsberg, R., Holzer-Popp, T., Paul, F., Sandven, S., Sathyendranath, S., van Roozendael, M., Wagner, W.: The esa climate change initiative: Satellite data records for essential climate variables". *Bulletin of the American Meteorological Society* **94**(10), 1541 – 1552 (2013)
14. Huang, X., Wang, C., Purushotham, S., Wang, J.: Vdam: Vae based domain adaptation for cloud property retrieval from multi-satellite data. In: Proceedings of the 30th International Conference on Advances in Geographic Information Systems. ACM SIGSPATIAL (2022)

15. Ilteralp, M., Ariman, S., Aptoula, E.: A deep multitask semisupervised learning approach for chlorophyll-a retrieval from remote sensing images. *Remote Sensing* **14**(1), 18 (2021)
16. Kuga, R., Kanezaki, A., Samejima, M., Sugano, Y., Matsushita, Y.: Multi-task learning using multi-modal encoder-decoder networks with shared skip connections. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. pp. 403–411 (2017)
17. Liu, C., Yang, S., Di, D., Yang, Y., Zhou, C., Hu, X., Sohn, B.J.: A machine learning-based cloud detection algorithm for the himawari-8 spectral image. *Advances in Atmospheric Sciences* **39**(12), 1994–2007 (2022)
18. Miranda, F.M., Köhnecke, N., Renard, B.Y.: Hiclass: a python library for local hierarchical classification compatible with scikit-learn. *Journal of Machine Learning Research* **24**(29), 1–17 (2023)
19. Platnick, S., Meyer, K., Wind, G., Holz, R.E., Amarasinghe, N., Hubanks, P.A., Marchant, B., Dutcher, S., Veglio, P.: The nasa modis-viirs continuity cloud optical properties products. *Remote sensing* **13**(1), 2 (2020)
20. Sayer, A.M., Lelli, L., Cairns, B., van Diedenhoven, B., Ibrahim, A., Knobelspiesse, K.D., Korkin, S., Werdell, P.J.: The chroma cloud-top pressure retrieval algorithm for the plankton, aerosol, cloud, ocean ecosystem (pace) satellite mission. *Atmospheric Measurement Techniques* **16**(4), 969–996 (2023)
21. Shao, Z., Pan, Y., Diao, C., Cai, J.: Cloud detection in remote sensing images based on multiscale features-convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing* **57**(6), 4062–4076 (2019)
22. Toshniwal, S., Tang, H., Lu, L., Livescu, K.: Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition. *arXiv preprint arXiv:1704.01631* (2017)
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
24. Voigt, A., Albern, N., Ceppi, P., Grise, K., Li, Y., Medeiros, B.: Clouds, radiation, and atmospheric circulation in the present-day climate and under climate change. *Wiley Interdisciplinary Reviews: Climate Change* **12**(2), e694 (2021)
25. Wang, C., Platnick, S., Meyer, K., Zhang, Z., Zhou, Y.: A machine-learning-based cloud detection and thermodynamic-phase classification algorithm using passive spectral observations. *Atmospheric Measurement Techniques* **13**(5), 2257–2277 (2020)
26. Wang, Q., Zhou, C., Zhuge, X., Liu, C., Weng, F., Wang, M.: Retrieval of cloud properties from thermal infrared radiometry using convolutional neural network. *Remote Sensing of Environment* **278**, 113079 (2022)
27. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7794–7803 (2018)
28. Werdell, P.J., Behrenfeld, M.J., Bontempi, P.S., Boss, E., Cairns, B., Davis, G.T., Franz, B.A., Gliese, U.B., Gorman, E.T., Hasekamp, O., Knobelspiesse, K.D., Manzano, A., Martins, J.V., McClain, C.R., Meister, G., Remer, L.A.: The plankton, aerosol, cloud, ocean ecosystem mission: Status, science, advances. *Bulletin of the American Meteorological Society* **100**(9), 1775 – 1794 (2019)
29. Yang, Y., Sun, W., Chi, Y., Yan, X., Fan, H., Yang, X., Ma, Z., Wang, Q., Zhao, C.: Machine learning-based retrieval of day and night cloud macrophysical parameters over east asia using himawari-8 data. *Remote Sensing of Environment* **273**, 112971 (2022)