

Using Unsupervised and Supervised Learning and Digital Twin for Deep Convective Ice Storm Classification

Jason Swope¹, Steve Chien¹, Emily Dunkel¹, Xavier Bosch-Lluis¹, Qing Yue¹ and William Deal²

¹Jet Propulsion Laboratory, California Institute of Technology

²Northrop Grumman

jason.swope@jpl.nasa.gov

Abstract

Smart Ice Cloud Sensing (SMICES) is a small-sat concept in which a primary radar intelligently targets ice storms based on information collected by a lookahead radiometer. Critical to the intelligent targeting is accurate identification of storm/cloud types from eight bands of radiance collected by the radiometer. The cloud types of interest are: clear sky, thin cirrus, cirrus, rainy anvil, and convection core.

We describe multi-step use of Machine Learning and Digital Twin of the Earth’s atmosphere to derive such a classifier. First, a digital twin of Earth’s atmosphere called a Weather Research Forecast (WRF) is used generate simulated look-ahead radiometer data as well as deeper “science” hidden variables. The datasets simulate a tropical region over the Caribbean and a non-tropical region over the Atlantic coast of the United States. A K-means clustering over the scientific hidden variables was utilized by human experts to generate an automatic labelling of the data - mapping each physical data point to cloud types by scientists informed by mean/centroids of hidden variables of the clusters. Next, classifiers were trained with the inputs of the simulated radiometer data and its corresponding label. The classifiers of a random decision forest (RDF), support vector machine (SVM), Gaussian naïve bayes, feed forward artificial neural network (ANN), and a convolutional neural network (CNN) were trained. Over the tropical dataset, the best performing classifier was able to identify non-storm and storm clouds with over 80% accuracy in each class for a held-out test set. Over the non-tropical dataset, the best performing classifier was able to classify non-storm clouds with over 90% accuracy and storm clouds with over 40% accuracy. Additionally both sets of classifiers were shown to be resilient to instrument noise.

1 Introduction

High altitude ice clouds, covering more than 50% of the Earth’s surface are often produced from high-impact deep convection events [Luo and Rossow, 2004], and are strong modulators of Earth’s weather and climate [Stephens, 2005][Bony *et al.*, 2006]. High altitude ice clouds play a significant role in the Earth’s energy balance and hydrologic cycle through their effects on radiative feedback and precipitation, and are therefore crucial for life on Earth. The SMICES mission is designed to dramatically increase our knowledge of these deep convective ice storms.

Most satellites blindly image nadir (directly below the orbiting spacecraft) without any knowledge of the phenomena they are observing. Even though global cloud coverage is roughly 2/3 of the Earth [King *et al.*, 2013], deep convective storms are far rarer than all clouds, representing only a fraction of a percent of all data. Additionally, radar used to study ice storms has a small footprint of $4km \times 4km$. This means that if the radar were to blindly image nadir, very few storms would be captured. Additionally, because radar uses a lot of energy, SMICES will only be able to capture data for about 20% of the time. These constraints lead to the need for active targeting, which will allow the mission to maximize its scientific return by selecting which clouds are analyzed.

This paper focuses on the classification system of the SMICES mission. We describe the use of unsupervised machine learning, digital twin (WRF), and supervised machine learning to address the challenge of developing a radiometer-based classifier for deep convective ice storms.

The remainder of this paper is organized as follows. In Section 3 we provide context on the problem and an overview of the classifiers and data that will be used. Section 4 describes the tropical and non-tropical datasets and how an automatic labelling system was established. Section 5 and 6 explain the different classifiers utilized and how they were set up for this experiment. Section 7 reviews the results of each classifier on both datasets. Section 8 explores the impact of noise on the accuracy of the classifiers. Section 9 and 10 discuss the results of the classifiers in the context of overall performance and outlines future work. Finally, in section 11 we summarize our conclusions.

2 Related Work

There are examples of work on cloud detection in other use cases such as cloud avoidance. On the EMIT mission cloud detection is used on the hyperspectral data to screen out covered sections of images [Oaida *et al.*, 2022]. Other uses of cloud screening can be found in onboard flood and cryospheric classification [Ip *et al.*, 2006][Doggett *et al.*, 2006]. Further work on screening clouds out of data data has been pursued to prepare for the unprecedented data volumes that will come from future missions [Thompson *et al.*, 2014].

In our work, we use supervised learning for storm classification at the pixel level. We focus on storm classification from satellite data instead of cloud detection, and we use simulated radiometer data from a satellite. In addition, we use a digital twin in combination with sparse labeling and unsupervised clustering to facilitate label generation.

Prior work has explored the use of unsupervised learning to cluster cloud data. Clustering has been explored using three-dimensional histograms applied to multi-spectral satellite imagery, using the visible, IR, and water-vapor channels [Desbois *et al.*, 1982]. Later work improved upon this by introducing textural parameters and processing larger datasets at different times [Seze and Desbois, 1987]. Another work attempted to reproduce the class clusters using Probabilistic Self-organizing Maps [Ambroise *et al.*, 2000].

Our clustering approach differs from the above work, in that we use simulated science parameters from a digital twin to do clustering, as opposed to using features derived from imagery. To the best of our knowledge, there is currently no other work with digital twins being used for the purpose of facilitating label generation in cloud imagery.

3 Background

The SMICES classification problem is to correctly classify different cloud types utilizing the information that the onboard radiometer collects. The classifier needs to be able to correctly identify the cloud as one of five cloud types: clear sky, thin cirrus, cirrus, rainy anvil, and convection core. The scientifically highest value cloud is the convection core followed by the rainy anvil cloud. With respect to the SMICES targeting algorithms, only the rainy anvil and convection core clouds are actively targeted. Therefore, it is extremely important for the classifiers to be able to distinguish between the rainy anvil, the convection core and the three non-storm cloud types. An accurate classification of these cloud types has been shown to enable a gain of capturing convection core clouds by a factor of 24 and rainy anvil clouds by a factor of 2 [Swope *et al.*, 2021].

We first use clustering and human experts to label the digital twin / WRF data. The digital twin datasets include three scientific variables of ice water path (IWP), median particle size (μm), and median cloud top height (m), as well as the eight bands of radiance previously mentioned. These three scientific variables are used to automatically generate labels for the data with the help of scientists as explained in section 4.2. This avoids the need to manually label the data.

We then use these labels with the digital twin (simulated) radiometer data to train the classifiers. During testing, the

classifiers will only have access to the radiance values since those are the only data available in an operational setting.

A set of classifiers were trained and tested on two separate regional datasets: a tropical dataset over the Caribbean, and a non-tropical dataset of the Atlantic Coast of the United States. The physics of a given cloud type can differ depending on whether it is in a tropical or non-tropical region. For example, there is ice in the deep convective core in non-tropical regions but not in tropical regions, and this affects the radiance values. Thus, we train separate classifiers for each dataset. These classifiers include a random decision forest (RDF), support vector machine (SVM), Gaussian Naïve Bayes, feed forward artificial neural network (ANN), and convolutional neural network (CNN). In an operational setting, we would select the model that corresponds to the geographical region the satellite is over.

4 Data

4.1 Datasets

Two regional datasets were used in this work: a tropical and non-tropical dataset. Both datasets were created through the Global Weather Research and Forecasting (GWRf) model [Skamarock *et al.*, 2019]. The GWRf is a state-of-the-art physics-based weather model. It is used to create computationally expensive datasets that we use as a digital twin to real climate data. In our case the model generated the brightness temperatures for different cloud types along the bands of Tb250+0.0, Tb310+2.5, Tb380-0.8, Tb380-1.8, Tb380-3.3, Tb380-6.2, Tb380-9.5, and Tb670+0.0, as well as the scientific variables of ice water path, median particle size, and median cloud top height.

The tropical dataset is located in the Caribbean. This dataset contains 13 images that are 119x208 pixels with a pixel size of 15km for a spatial extent of 1,785km x 3,120km. Each image is a snapshot of the same area in one-hour intervals.

The non-tropical dataset is located in the Atlantic Coast of the United States. The data contains 29 image cutouts that are 1998x270 pixels with a pixel size of 1.33km for a spatial extent of 2,657km x 359km. The dataset combines to form three images over the same area in 12-hour time intervals. Each larger image is constructed of 10 image cutouts stacked vertically. The total size of a full image is 1998x2700 pixels with a spatial extent of 2,657km x 3,591km. The last cutout of the third time interval was incomplete and therefore left out of this study. Therefore, we only have 29 image cutouts.

The same cloud types in the tropical vs. non-tropical datasets do not necessarily correspond to similar scientific features or radiance values. Because of this, as well as the large difference in resolution (pixel size of 15km vs. 1km), we treat the tropical and non-tropical datasets independently in this work.

4.2 Data Labelling Using Clustering and Digital Twin

Manually labelling our data was not a feasible option given the extremely large number of pixels, and the difficulty in identifying the cloud type based only on radiance values. We




Figure 1: Cluster values mapped to their labels on the tropical dataset

used a digital twin of the radiometer to solve this problem. The digital twin helps in two ways. It gives us access to scientific variables that are not directly measurable in nature and also allows us to generate a large amount of storm data for training and evaluation.

The scientific variables of ice water path, median particle size, and median cloud top height are significant because they give us a way to map a pixel's values to a specific cloud type. Unlike the radiance values, the scientists are able to map these scientific variables to the five cloud types they identified. While this solves the mapping problem, an automated method needed to be developed due to the large number of pixels

The approach to automate the labelling is as follows. First, we cluster the data into representative clusters based on the scientific variables. Then we can map each cluster center to the cloud type that it corresponds to. Each cluster's mapping serves as the label for every pixel within that cluster.

K-means was utilized as our clustering method. This approach requires us to identify the proper number of clusters to accurately represent our dataset before mapping the clusters to the specific cloud types. A failure to accurately represent our data will cause a poor classification of pixels. The proper number of clusters is determined by analyzing if the number of clusters selected maximizes the within cluster coherence and the between cluster separation. A silhouette score was used to calculate the effectiveness of the number of clusters over the data.

A silhouette score, which ranges from -1 to 1, gives a measure of the cluster's fit. A score of 1 means the clusters are well distinguished, a score of 0 means the clusters are indifferent, and a score of -1 means that the clusters are incorrect. The equation used to obtain this score is: $((b-a)/\max(a,b))$ where a represents the intra-cluster distance (average distance between each point within a cluster) and b represents the average nearest cluster distance (average distance between the instances of the next closest cluster). Clustering was performed separately for the tropical and non-tropical datasets.

There are three local maximums at 5, 9, and 16 clusters in the tropical dataset. It is possible to map multiple clusters from K-means to one cloud type if multiple clusters have

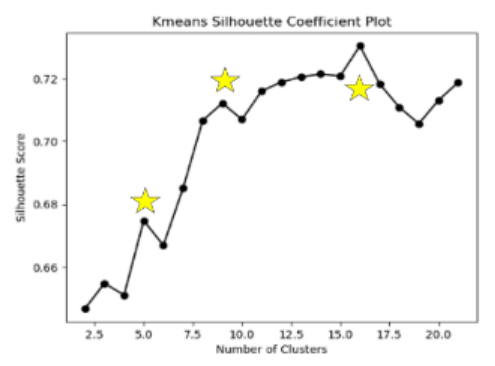


Figure 2: Silhouette scores over the tropical dataset

similar scientific values. The difference between the silhouette scores of the local maximum at five clusters (.675) is only .06 away from the silhouette score at 16 clusters (.735), which is an insignificant increase. We chose to use five clusters due to the high silhouette score and the easier mapping to the five alluded cloud types.

Silhouette scores were also calculated for a random sample taken from the non-tropical dataset. The scores for all of the clustering values were over .97, which suggests a very good clustering for the dataset. We chose to continue to use a 5 cluster K-means for the non-tropical dataset for its high silhouette score and easy mapping to the original cloud labels.

The scientists assigned each cluster to its respective label based on the mean and standard deviation of its centroid. Overall, the clusters gradually increased in IWP and median particle size while decreasing in median cloud top height. The only exception to this rule was the cluster with all values set to 0, which we assign to the clear sky class. This relationship between the clusters was used by the scientists to identify the labels. In general, a higher IWP and median particle size while having lower median cloud top height correlates to a stronger storm cloud. These mappings for the tropical dataset are demonstrated in figure 1.

5 Classifiers

The ground truth values for each cloud class are taken to be the labels assigned through the clustering method described above. Now that we have labels, we can train the classifiers to predict the cloud class using only data available in orbit: the radiometer data. During training, the classifiers see both the radiometer data and the class labels, and at test time, see only the radiometer data. The classifiers operate at the single pixel level; given a radiometer measurement at a given pixel, what is the cloud class at that pixel. Thus, our current classifiers do not take into account neighboring pixel values. Classifier performance was evaluated using 5-fold cross-validation, and testing was performed on a separate held-out test set.

For this study we explored the following classifiers: random decision forest (RDF), support vector machine (SVM), Gaussian Naïve Bayesian, a feed forward artificial neural network (ANN), and a convolutional neural network (CNN).

An important factor in the classifier configuration is how balanced the different classes within the datasets are. The classes are heavily skewed away from the most important cloud type, the convection core, in both the tropical and non-tropical datasets.

	Clear Sky	Thin Cirrus	Cirrus	Rainy Anvil	Convection Core
Full Dataset	23.8%	14.0%	29.2%	32.0%	1.0%
Train Dataset	24.1%	13.6%	30.0%	31.3%	1.0%
Test Dataset	22.9%	15.7%	26.3%	34.2%	0.9%

(a) Distributions over the Tropical Dataset

	Clear Sky	Thin Cirrus	Cirrus	Rainy Anvil	Convection Core
Full Dataset	60.2%	10.4%	9.6%	15.9%	3.8%
Train Dataset	59.3%	8.2%	11.1%	17.5%	3.8%
Test Dataset	62.2%	15.4%	6.2%	12.4%	3.8%

(b) Distributions over the Non-Tropical Dataset

Table 1: Cloud type distribution over the tropical and non-tropical datasets and their respective train and test datasets

Table 1 displays the distribution of cloud type in the tropical and non-tropical dataset. The two most important classes, convection core and rainy anvil, only make up around 32% of the dataset. Therefore, if we only classified every pixel as clear sky, thin cirrus, or cirrus we would be able to attain a 68% accuracy, although we would miss the important classes. The class distribution of the non-tropical dataset (table 1) is significantly more skewed towards clear than the tropical dataset. About 80% of this dataset is non-storm clouds. To help the classifiers overcome these unbalanced datasets, the random forest and support vector machine classifiers are trained with weights adjusted for the class imbalance. When it is noted that the weights are equalized, we are stating that the weights have been rebalanced so that every class is weighted equally.

6 Experimental Design

Separate classifiers were trained for tropical and non-tropical regions, using the data corresponding to that region. For the tropical dataset the first ten images were used as a training and validation set. The remaining three images then served as the test set. The test and validation set for the non-tropical dataset was created from the first two time steps. These two timesteps are comprised of the first 20 image cutouts in the dataset. The remaining 9 image cutouts that makeup the final image was used as the test set. The distribution of the cloud classes in each training and test dataset is shown in table 1. We did not build classifiers using a combined dataset due to the different pixel sizes of each dataset and physical differences found in tropical and non-tropical clouds.

The classifiers are evaluated on their performance by analyzing their accuracy over the three cloud classes of non-storm (clear, thin cirrus, and cirrus), rainy anvil, and convection core. The classifiers were trained in the following method:

1. Trained on the original five-class labelling created through K-means. These classifiers output a five-class labelling of the clouds which is converted into the three-class accuracy¹

The performance of each classifier will be presented over the three-class problem of identifying non-storm clouds (clear, thin cirrus, or cirrus), rainy anvil clouds, and convection core clouds. The simplification from five to three classes is due to SMICES only actively targeting rainy anvil and convection core clouds, with a preference for convection core. Therefore, the meaningful distinctions are between those three classes. However, not all misclassifications are equal. A misclassification of a non-storm cloud as rainy anvil or convection core cloud could trigger the use of radar power for a non-desired target. This would be more costly than a misclassification between a convection core and rainy anvil cloud since the radar would still be collecting scientifically significant data. The performance will also be discussed in the context to the two class problem of non-storm clouds (clear, thin cirrus, and cirrus), and storm clouds (rainy anvil and convection core). This analysis highlights how many significant misclassifications are being made by each classifier.

7 Classifier Results

7.1 Random Decision Forest (RDF)

We use scikit-learn’s implementation of the Random Forest Classifier in this work [Pedregosa, 2011]. We found that 32 trees and a maximum depth of 14 was optimal based on the training and validation sets.

One challenge with the storm datasets is their unbalanced nature. Storms are quite rare in the sky, and the center of storms (convection core clouds) are even rarer. Therefore, the classifier is going to be more prone to classifying clouds as non-storm clouds, since that is the most dominant class. To

¹We also explored training the data on three class labeled data, however this produced similar results. This paper will focus on the 5 class results.

solve this, the weights have been adjusted in the RDF so that every class is weighted equally.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	44226 88.38%	5751 11.49%	63 0.13%
Rainy Anvil	3858 16.55%	19043 81.70%	408 1.75%
Convection Core	109 12.02%	584 64.39%	214 23.59%

Table 2: Confusion Matrix of the RDF classifier on the tropical dataset, max depth 14, number of trees 32, weights equalized by class, 3 class labeled data

Table 2 shows the RDF performance. On the tropical data the classifier was able to accurately classify non-storm clouds and rainy anvil clouds with 88% and 82% and accuracy respectively. Even though the classifier failed to accurately identify convection core clouds, the majority of convection core misclassifications were as rainy anvil clouds. When looking at the storm cloud accuracy the classifier was able to identify storm clouds with 84% accuracy.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	3539649 89.61%	246916 6.25%	163505 4.14%
Rainy Anvil	254077 46.25%	196388 35.75%	98939 18.01%
Convection Core	273330 76.85%	38958 10.95%	43378 12.20%

Table 3: Confusion Matrix of the RDF classifier on the non-tropical dataset, max depth 14, number of trees 32, weights equalized

The non-tropical dataset proved to be significantly more difficult to classify than the tropical dataset. The accuracy for the storm cloud classes (rainy anvil and convection core) decreased by about 2x. The non-storm class accuracy remained relatively unchanged when compared to the tropical dataset. When looking at the non-storm/storm cloud accuracy the RDF achieved 90% non-storm accuracy and 42% storm accuracy.

7.2 Support Vector Machine (SVM)

The support vector machine used was the linear support vector classifier (linear SVC) from Scikit Learn. This classifier is suggested if the dataset contains greater than tens of thousands of data points by Scikit Learn. The linear SVC is effectively an support vector machine with a linear kernel.

Table 4 demonstrates the accuracy of the SVM classifier with a linear kernel on the tropical dataset. The classifier accurately classifies the non-storm and rainy anvil classes with 82% and 81% accuracy. The SVM failed to accurately classify the convection core class. Similar to the RDF, most of its convection core misclassifications were labeled as rainy anvil. The storm cloud accuracy for this classifier is 68%.

The SVM classifier performed worse in the storm classes on the non-tropical dataset than it did on the tropical dataset.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	44821 82.22%	9552 17.52%	139 0.25%
Rainy Anvil	3199 16.87%	15418 81.31%	344 1.81%
Convection Core	173 22.09%	408 52.11%	202 25.80%

Table 4: Confusion Matrix of the linear SVC classifier on the tropical dataset, weights balanced

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	3464133 89.15%	307787 7.92%	113775 2.93%
Rainy Anvil	403130 57.14%	253696 35.96%	48740 6.91%
Convection Core	199433 75.58%	40114 15.20%	24332 9.22%

Table 5: Confusion Matrix of the linear SVC classifier on the non-tropical dataset, weights balanced

While the non-storm classification accuracy increased to 89%, the accuracies in both storm classes decreased by more than half when compared to the tropical dataset performance. The storm class accuracy over the non-tropical dataset 40%.

7.3 Gaussian Naïve Bayes

The Gaussian Naïve Bayes classifier used is from Scikit learn. The likelihood of each feature is assumed to be Gaussian.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	47188 74.76%	15715 24.90%	217 0.34%
Rainy Anvil	736 7.45%	8926 90.33%	220 2.23%
Convection Core	269 21.45%	737 58.77%	248 19.78%

Table 6: Confusion Matrix of the Gaussian Naïve Bayes Classifier on the tropical dataset

Table 6 contains the results of the Gaussian Naïve Bayes classifier when trained on the tropical dataset. It achieved the best classification in the rainy anvil class with 90%, however struggled with the non-storm and convection core classes compared to the other classifiers. Most of the convection core misclassifications were made as rainy anvil classifications. The two class accuracy over the tropical dataset is 75% for non-storm and 42% for storm.

Table 7 shows the performance of the GNB classifiers over the non-tropical dataset. Similar to the RDF and SVM, the GNB struggled to accurately classify the separate classes. While the non-storm accuracy increased when compared to the tropical dataset, its accuracy in both storm classes decreased by more than half. The two class accuracy over the tropical dataset is 92% for non-storm and 55% for storm.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	3393184 92.03%	187456 5.08%	106523 2.89%
Rainy Anvil	599488 56.15%	394857 36.98%	73349 6.87%
Convection Core	74021 73.81%	19290 19.24%	6972 6.95%

Table 7: Confusion Matrix of the Gaussian Naïve Bayes Classifier on the non-tropical dataset

7.4 Neural Networks

We also applied the use of simple neural networks using Keras [Chollet and others, 2015] and TensorFlow [Abadi *et al.*, 2015].

Feed Forward Artificial Neural Network (ANN)

We pass each 8-band pixel into a feed forward ANN with 2 hidden layers, 32 nodes in each hidden layer, a dropout rate of 0.1, and a softmax activation at the final layer. We used ADAM [Kingma and Ba, 2017] as our optimizer. Model performance was not affected by small changes in these parameters. The model output is a vector with probability of each cloud-type class, and we take the Argmax to get the predicted class.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	46202 95.87%	1889 3.92%	102 0.21%
Rainy Anvil	13044 51.40%	12043 47.46%	290 1.14%
Convection Core	148 21.57%	389 56.71%	149 21.72%

Table 8: Confusion Matrix of the ANN on the tropical dataset

The ANN performance over the tropical dataset was relatively poor compared to the past classifiers. While the classifier achieved around 95% accuracy in the non-storm class, it struggled to identify both the convection core and rainy anvil clouds (table 8). Even in the two class problem neither classifier was able to identify storm clouds with 60% accuracy.

The ANN was not effective on the non-tropical dataset, and labelled all pixels as the non-storm class during training and validation. This is probably due to the larger skew of the non-tropical dataset towards the non-storm class.

Convolutional Neural Network (CNN)

We also test a single-pixel CNN model. Our architecture uses 3 stacked 1-d convolutional layers, where the convolutions are applied over the bands. Due to the small input size, we did not apply pooling. 6, 12, and 24 filters were used for each convolutional layer respectively, and a rectified linear unit (relu) activation was used. We found a convolutional filter shape of 1x3 to be optimal. The last layers of our net included two fully connected layers with 32 nodes, a dropout of 0.1, and a final softmax activation. ADAM optimizer was used to train the network.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	43893 91.08%	4190 8.69%	110 0.23%
Rainy Anvil	5926 23.35%	19000 74.88%	449 1.77%
Convection Core	83 12.06%	405 58.87%	200 29.07%

Table 9: Confusion Matrix of the CNN on the tropical dataset

Table 9 show the CNN results on the tropical dataset. The classifier was able to classify both non-storm and rainy anvil clouds accurately with over 91% accuracy in non-storms and 75% accuracy in rainy anvil. The classifier still struggled with convection core. It was able to achieve a storm cloud accuracy of 83% which is very close to the RDF performance.

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Clear, Thin Cirrus, and Cirrus	3348548 90.82%	307505 8.34%	30902 0.84%
Rainy Anvil	539227 59.88%	322980 35.86%	38341 4.26%
Convection Core	228529 85.39%	30029 11.22%	9079 3.39%

Table 10: Confusion Matrix of the CNN on the non-tropical dataset

The CNN experiences a similar outcome to the ANN when trained on the non-tropical dataset. The heavier skew towards non-storm clouds caused the classifier to classify the majority of the test dataset as non-storm. The classifier achieved a 91% accuracy in non-storm clouds and a 44% accuracy in storm clouds. For future work, we plan to upsample the minority classes when training our neural networks.

8 Noise

In flight we expect the sensor to generate Gaussian noise that is independent along each radiance band. The Tb380 band is expected to have around 5 kelvin of noise, and the rest of the bands are expected to have 1 kelvin.

To test the impact of the expected noise on the classifiers we generate random values from a Gaussian distribution with a mean of 0 and a standard deviation of the expected noise for each band. These values are then independently added to each band to create a noisy dataset. The impact of noise is determined by the total decrease in accuracy of the three-class RDF classifier when tested on the clean held-out test set, and on the same testset, but with noise applied.

1. Tropical dataset: RDF classifier with a max depth of 14, 32 trees, and weights equalized was trained on the first eight images of the tropical dataset. It was then tested on the remaining five images after expected noise was applied to those data values.
2. Non-tropical dataset: RDF classifier with a max depth of 14, 32 trees, and weights equalized was trained on the first image of the non-tropical dataset (10 image cutouts). It was then tested on the second image (10 image cutouts) after expected noise was applied noise was applied to those data values.

Over the tropical dataset the expected noise in flight decreased the accuracy of the RDF by 4%. When the noise was applied to the non-tropical dataset the accuracy of the RDF decreased by 2%. This demonstrates that the RDF classifier is relatively robust to noise in both datasets

9 Discussion

Multiple classifiers were able to accurately distinguish between the two-class problem in the tropical dataset. It is important to understand that even less accurate classifiers can be used to significantly increase yield of the SMICES mission concept. In Table 11a the first row shows the distribution of pixels acquired if sampling pixels at random (e.g. not using any classification). Row two shows the distribution if we sample pixels the RDF classifies as Rainy Anvil. Note that in the chart, the term "Sample Labelled" means what the classifier thinks is the correct class. We still get some non-storm and Convection Core pixels, since the classifier is not 100% accurate. Row three shows the distribution if we only sample pixels the RDF classifies as Convection Core. Again, we get other classes as well due to inaccuracies in the classifier. This data shows that sampling a Rainy Anvil or Convection Core classified pixel is far more fruitful scientifically than a random pixel.

Previous work by Swope [Swope *et al.*, 2021] has evaluated the improvements in mission return by running mission simulations using WRF datasets in which intelligent targeting attempts to preferentially target areas of Convection Core and Rainy Anvil but is limited by rarity of such pixels and mission energy and pointing constraints. In Table 11a, row 4 shows the distribution of pixels we would acquire if the classifier used in the simulation was 100% accurate. Row 5 shows the distribution of pixels (true labels) we actually acquire, quantifying how classification inaccuracy reduces the impact on return. However, this mode still dramatically outperforms uninformed (random) targeting. This highlights how preferential targeting is able to skew sampled pixels towards storm-cloud classes.

When looking at the non-tropical dataset, distinguishing between the two classes of Rainy Anvil and convection core is significantly harder. Again, even with the lower classification accuracy, Table 11b shows the expected pixels observed sampling from the Rainy Anvil and Convection Core labelled pixels compared to random sampling. Again, drawing on the mission simulations, the last two rows of table 11b show dramatically increased yields of Rainy Anvil and Convection Core measurements from intelligent targeting compared to uninformed (random) targeting even with imperfect classification.

10 Future Work

We would like to extend to additional regions beyond the Caribbean and Atlantic coast to make our classifiers more robust. A global dataset could be used to explore the possibility of a universal classifier that would work in any region (and season, and other conditions), however variations in atmospheric phenomena in different regions could make this

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Sample Randomly	67.0%	32.0%	1.0%
Sample Labelled as Rainy Anvil	16.6%	81.7%	1.8%
Sample Labelled as Convection Core	12.0%	64.4%	23.6%
Preferential Targeting (as labelled)	10.5%	64.2%	25.3%
Preferential Targeting (true labels)	23.2%	69.7%	7.1%

(a) Tropical Data

	Clear, Thin Cirrus, and Cirrus	Rainy Anvil	Convection Core
Sample Randomly	67.0%	32.0%	1.0%
Sample Labelled as Rainy Anvil	46.3%	35.8%	18.0%
Sample Labelled as Convection Core	76.9%	10.9%	12.2%
Preferential Targeting (as labelled)	44.6%	21.3%	34.1%
Preferential Targeting (true labels)	76.1%	14.1%	9.8%

(b) Non-Tropical Data

Table 11: Mission return impact of Intelligent Targeting with Classifier

difficult. Swapping between different regional classifiers in flight would be feasible.

The impact of the expected noise should also be analyzed more clearly. Even though the overall accuracy is not strongly impacted, it is important to know if any cloud types are being disproportionately affected by the noise or if is balanced.

Future work on the classifiers will expand beyond single pixel classification and take into account surrounding pixels. This should improve the accuracy because storm phenomena are not randomly distributed across the sky, instead they are clustered close together. Upsampling on the storm clouds in each dataset may also improve the overall performance, for the classifiers that did not have their weights equalized, due to how imbalanced the datasets are.

We intend to test these classifiers on real data from airborne tests of the SMICES radar.

11 Conclusion

We have described an effort to develop a classifier of deep convective storms based on radiometer data. Using a digital twin and K-means clustering we were able to generate labels for data. The results of the classifiers are promising for distinguishing deep convective storms in the tropical dataset. Further work still needs to be done for finer grained storm type discrimination. Additionally, identification of deep convective storms in non tropical data was more challenging. We also present results indicating that even moderate classification accuracies combined with intelligence instrument targeting are expected to enable significant improvements in mission return.

12 Acknowledgements

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [Abadi *et al.*, 2015] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro . . . , and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [Ambroise *et al.*, 2000] Christophe Ambroise, Geniève Sèze, Fouad Badran, and Sylvie Thiria. Hierarchical clustering of self-organizing maps for cloud classification. *Neurocomputing*, 30(1):47–52, 2000.
- [Bony *et al.*, 2006] S. Bony, R. Colman, V. M. Kattsov, R. P. Allan, C. S. Bretherton, J.-L. Dufresne, A. Hall, S. Hallegatte, M. M. Holland, W. Ingram, D. A. Randall, D. J. Soden, G. Tselioudis, and M. J. Webb. How well do we understand and evaluate climate change feedback processes? *J. Climate*, 19:3445–3482, 2006.
- [Chollet and others, 2015] Francois Chollet et al. Keras, 2015.
- [Desbois *et al.*, 1982] Michel Desbois, Genevieve Seze, and Gerard Szejwach. Automatic classification of clouds on meteosat imagery: Application to high-level clouds. *Journal of Applied Meteorology and Climatology*, 21(3):401–412, 1982.
- [Doggett *et al.*, 2006] T. Doggett, R. Greeley, S. Chien, R. Castano, B. Cichy, A.G. Davies, G. Rabideau, R. Sherwood, D. Tran, V. Baker, J. Dohm, and F. Ip. Autonomous detection of cryospheric change with hyperion on-board earth observing-1. *Remote Sensing of Environment*, 101(4):447–462, 2006.
- [Ip *et al.*, 2006] Felipe Ip, J.M. Dohm, V.R. Baker, T. Doggett, A.G. Davies, R. Castaño, S. Chien, B. Cichy, R. Greeley, R. Sherwood, D. Tran, and G. Rabideau. Flood detection and monitoring with the autonomous sciencecraft experiment onboard eo-1. *Remote Sensing of Environment*, 101(4):463–481, 2006.
- [King *et al.*, 2013] Michael D. King, Steven Platnick, W. Paul Menzel, Steven A. Ackerman, and Paul A. Hubanks. Spatial and temporal distribution of clouds observed by modis onboard the terra and aqua satellites. *IEEE Transactions on Geoscience and Remote Sensing*, 51(7):3826–3852, 2013.
- [Kingma and Ba, 2017] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [Luo and Rossow, 2004] Zhengzhao Luo and William B. Rossow. Characterizing tropical cirrus life cycle, evolution, and interaction with upper-tropospheric water vapor using lagrangian trajectory analysis of satellite observations. *Journal of Climate*, 17(23):4541 – 4563, 2004.
- [Oaida *et al.*, 2022] Bogdan V. Oaida, Amruta Yelamanchili, Christopher Wells, David R. Thompson, Robert O. Green, Matthew W. Bennett, Didier Keymeulen, Thang Pham, Daniel P. Poe, Lena Siskind, and Charlene Ung. Mapping earth’s dust-emitting regions from the iss with the emit imaging spectrometer. In *2022 IEEE Aerospace Conference (AERO)*, pages 1–12, 2022.
- [Pedregosa, 2011] F. et al Pedregosa. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [Seze and Desbois, 1987] Genevieve Seze and Michel Desbois. Cloud cover analysis from satellite imagery using spatial and temporal characteristics of the data. *Journal of Applied Meteorology and Climatology*, 26(2):287–303, 1987.
- [Skamarock *et al.*, 2019] C. Skamarock, Bogumiła Klemp, Jimmy Dudhia, O. Gill, Zhiqian Liu, Judith Berner, Wei Wang, G. Powers, Greg Duda, Dale M. Barker, and Xianguo Huang. A description of the advanced research wrf model version 4.1. 2019.
- [Stephens, 2005] Graeme L. Stephens. Cloud feedbacks in the climate system: A critical review. *J. Climate*, 18(2):237 – 273, 2005.
- [Swope *et al.*, 2021] Jason Swope, Steve Chien, Xavier Bosch-Lluis, Qing Yue, Peyman Tavallali, Mehmet Ogut, Isaac Ramos, Pekka Kangaslahti, William Deal, and Caitlyn Cooke. Using intelligent targeting to increase the science return of a smart ice storm hunting radar. July 2021.
- [Thompson *et al.*, 2014] David R. Thompson, Robert O. Green, Didier Keymeulen, Sarah K. Lundeen, Yasha Mouradi, Daniel Cahn Nunes, Rebecca Castaño, and Steve A. Chien. Rapid spectral cloud screening onboard aircraft and spacecraft. *IEEE Transactions on Geoscience and Remote Sensing*, 52(11):6779–6792, 2014.