DELTA: An Open-Source Framework to Simplify Machine Learning with Satellite Imagery

Michael von Pohle[1], Brian Coltin[2], and Scott McMichael[2]

[1]NASA Ames Research Center [USRA]
[2]NASA Ames Research Center

November 26, 2022

Abstract

DELTA (Deep Earth Learning, Tools, and Analysis) is an open-source framework developed at NASA to simplify running and training machine learning (ML) models on satellite imagery. Users new to machine learning can run existing ML models on satellite imagery with minimal setup and configuration. For experienced ML users, DELTA helps simplify data engineering, preprocessing steps, and reduces the need for boilerplate code that needs written to make satellite imagery datasets palatable for machine learning. This lets data scientists focus on model development while DELTA handles the imagery manipulation. This presentation will demonstrate DELTA's functionality and share some examples from an active project using it for flood mapping using imagery from multiple satellite sources.

```
train:
  loss: sparse_categorical_crossentropy
  metrics:
    - sparse_categorical_accuracy
  batch_size: 10
  epochs: 10
  validation:
    from_training: false
    images:
      type: tiff
      extension: _data.tif
      directory: validate
    labels:
      extension: _mask.png
      type: tiff
      directory: labels

mlflow:
  experiment_name: Landsat8 Clouds Example
```

```
dataset:
  images:
    type: tiff
    extension: _data.tif
    directory: train
  labels:
    extension: _mask.png
    type: tiff
    directory: labels
  classes:
    - 0:
        name: Shadow
        color: 0x000000
    - 1:
        name: Shadow over Water
        color: 0x000080
    - 2:
        name: Water
        color: 0x0000FF
    - 3:
        name: Snow
        color: 0x00FFFF
    - 4:
        name: Land
        color: 0x808080
    - 5:
        name: Cloud
        color: 0xFFFFFF
    - 6:
        name: Flooded
        color: 0x808000

io:
  tile_size: [512, 512] # is default size, DELTA has well chosen defaults for many of these
parameters
```

```yaml
dataset:
  images:
    type: tiff
    extension: _data.tif
    directory: train
  labels:
    extension: _mask.png
    type: tiff
    directory: labels
    preprocess:
      - substitute:
          mapping:
            - 0
            - 0
            - 2
            - 0
            - 0
            - 0
            - 0
  classes:
    - 0:
        name: Shadow
        color: 0x000000
    - 1:
        name: Shadow over Water
        color: 0x000080
    - 2:
        name: Water
        color: 0x0000FF
    - 3:
        name: Snow
        color: 0x00FFFF
    - 4:
        name: Land
        color: 0x808080
    - 5:
        name: Cloud
        color: 0xFFFFFF
    - 6:
        name: Flooded
        color: 0x808000

io:
  tile_size: [512, 512]
```

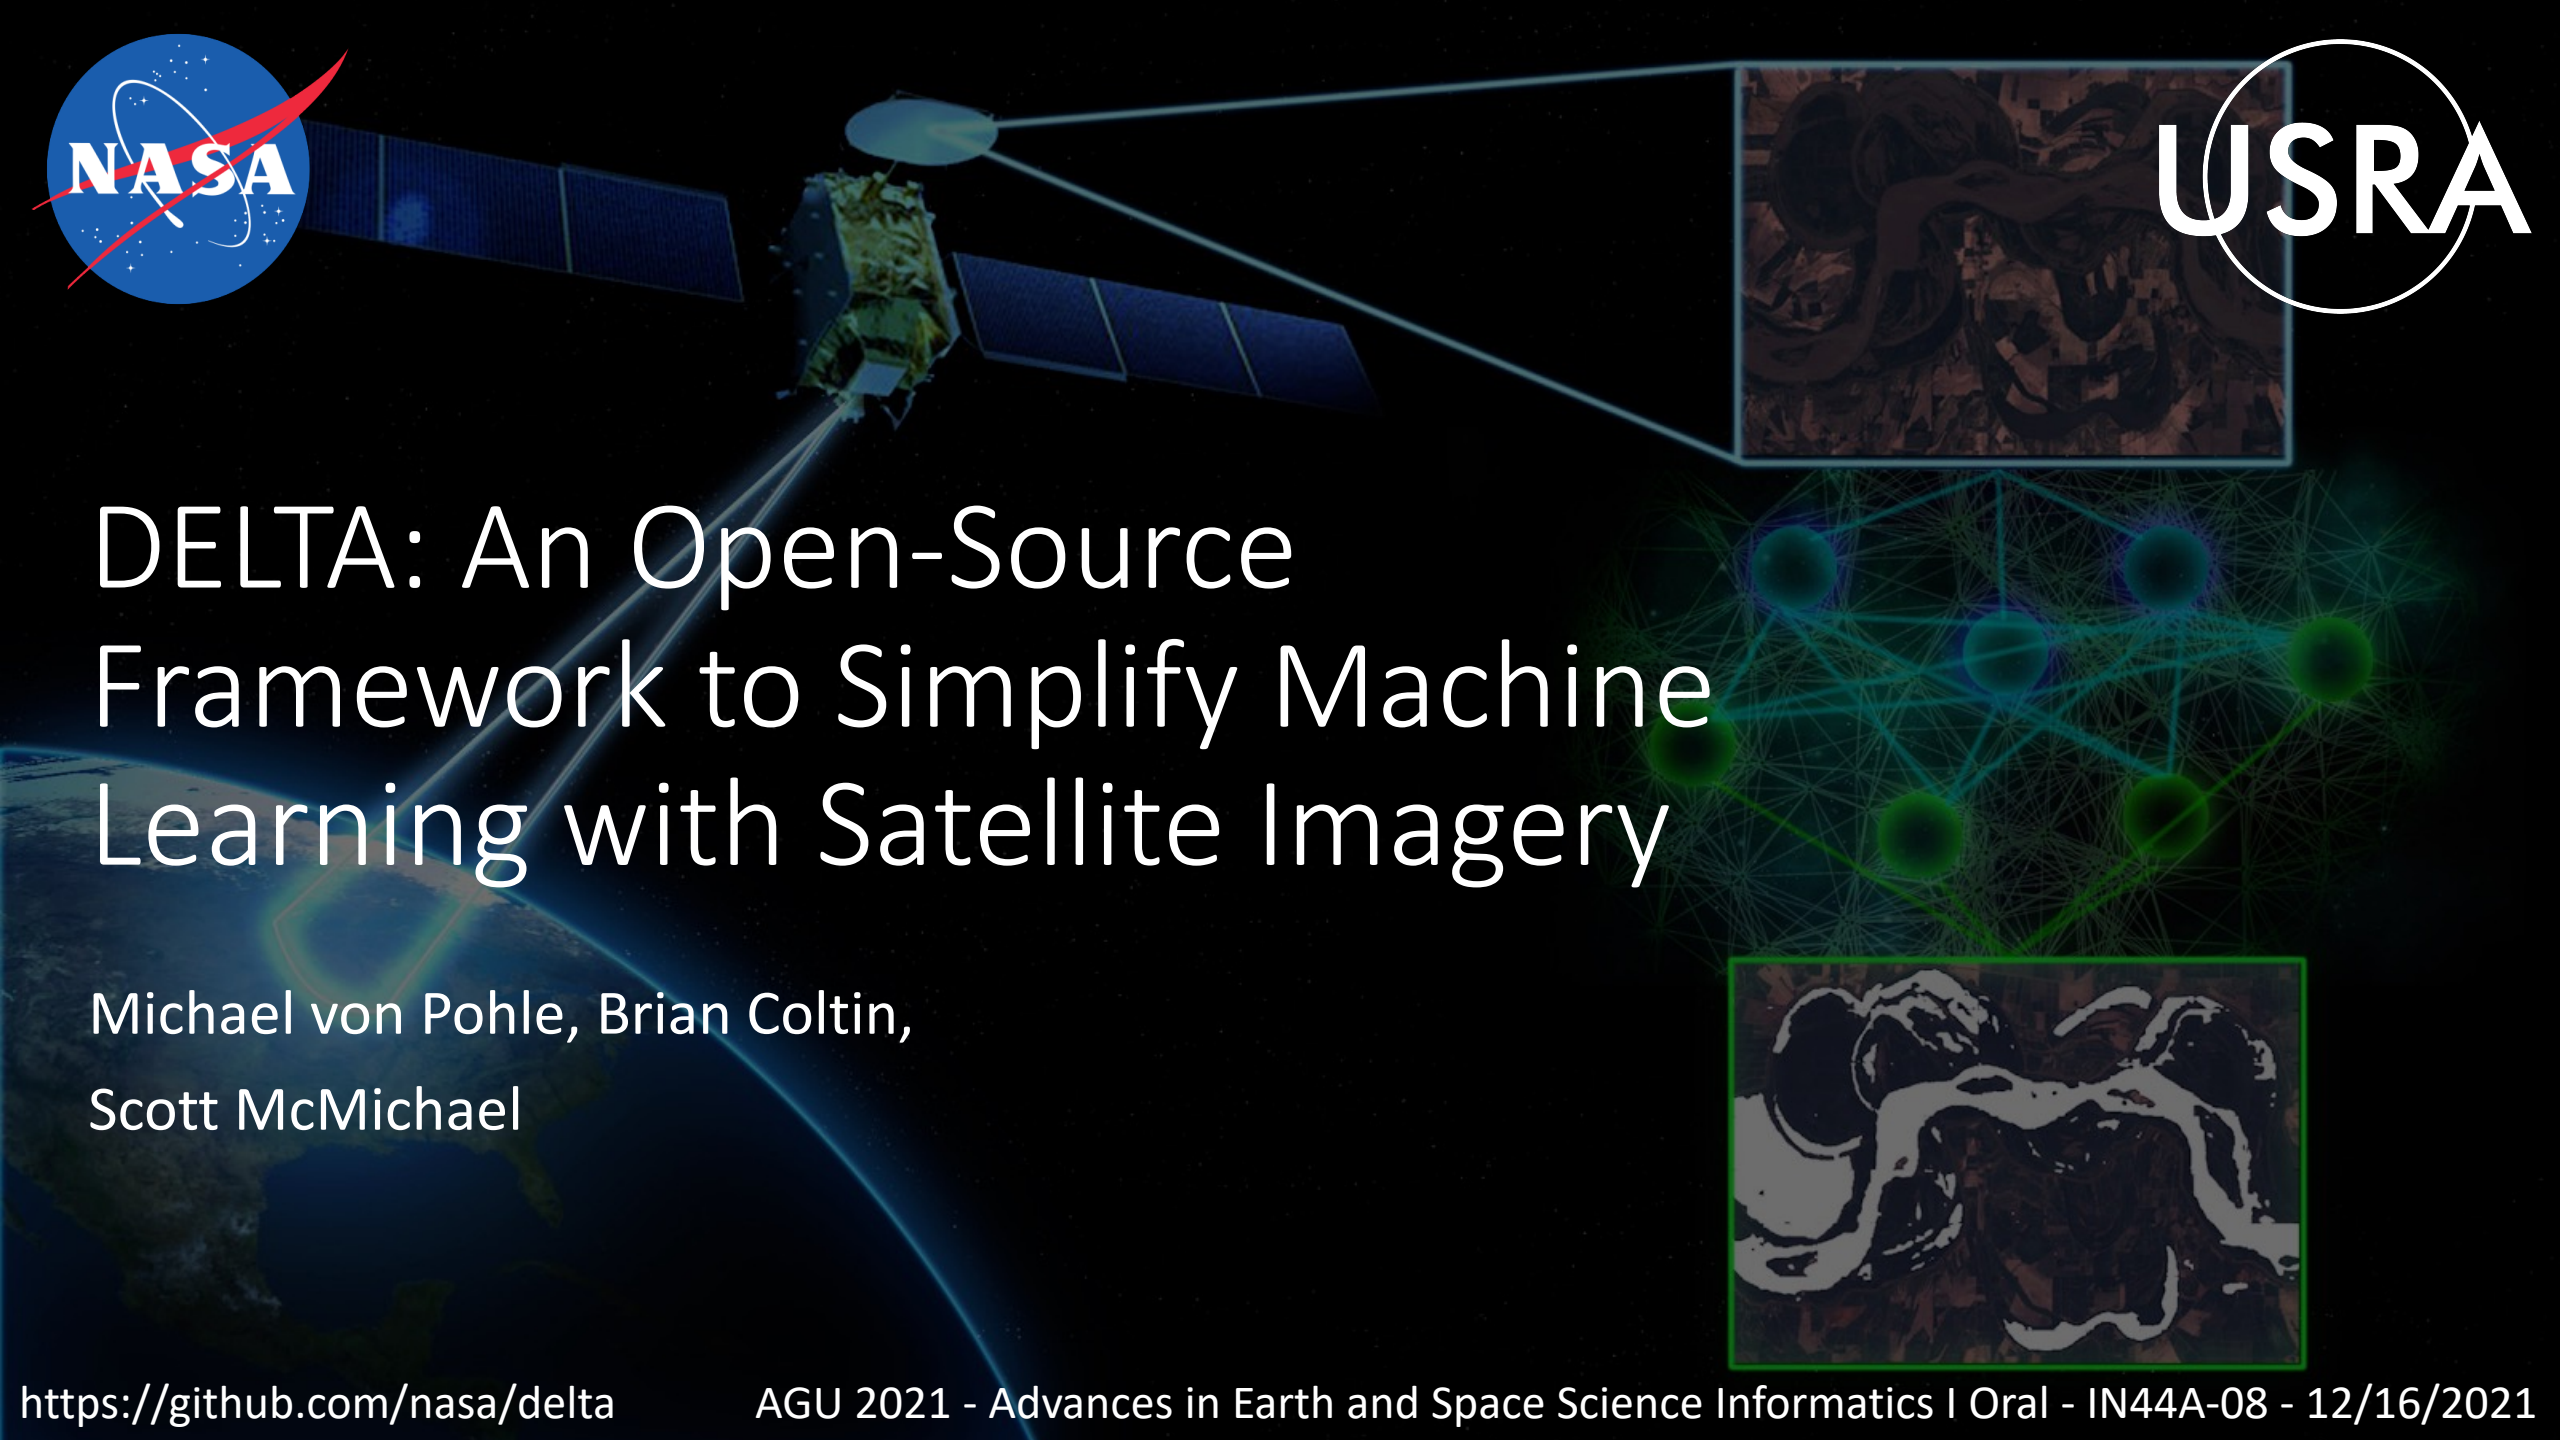

DELTA: An Open-Source Framework to Simplify Machine Learning with Satellite Imagery

Michael von Pohle, Brian Coltin,

Scott McMichael

https://github.com/nasa/delta

AGU 2021 - Advances in Earth and Space Science Informatics I Oral - IN44A-08 - 12/16/2021

DELTA Helps Address Two Big Problems

- Labor and time intensive preprocessing, postprocessing, and boilerplate code to turn satellite imagery into an ML-ready dataset
- Difficulty in applying trained models to satellite imagery. Especially in the earth science community among those with limited ML experience

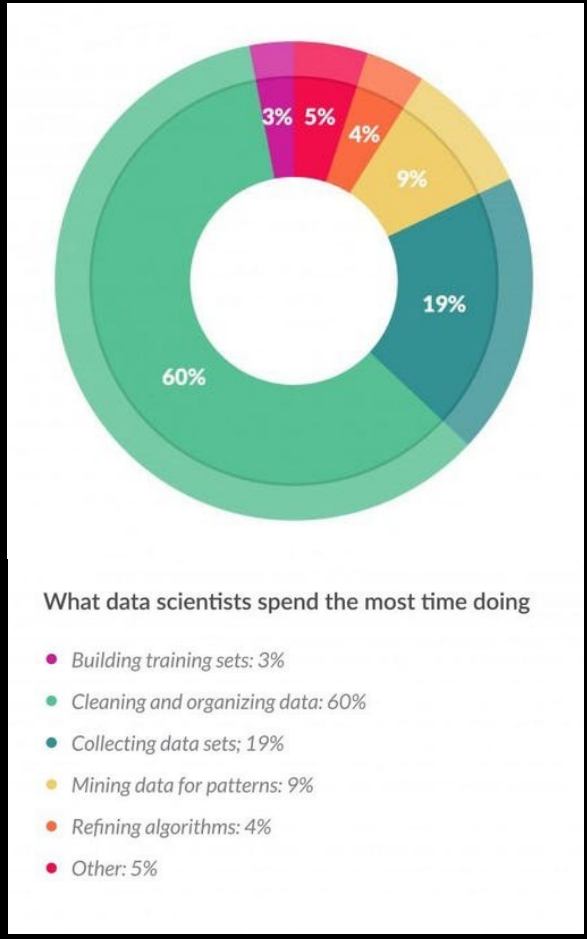

https://www.forbes.com/sites/gilpress/2016/03/23/data-preparation-most-time-consuming-least-enjoyable-data-science-task-survey-says/?sh=4a745d706f63

A Few Notes on Satellite Imagery

- Single images are often 1GB+ - 10k x 10x pixels or more
- A dataset of imagery can easily be several hundred GBs in size
- Often hyperspectral – more than RGB
- Can include metadata layers alongside data layers
- May not be rectangular

DELTA Features

- Works with any geospatial imagery/data that GDAL supports
- Has support for unpacking and preprocessing a growing list of common raw satellite imagery sources (Worldview, Sentinel, Landsat)
- Has a collection of common image preprocessing functions
- Performs tiling of imagery, training/inference, and restitching of predictions with adjustable overlap
- Works with h5 and SavedModel model formats
- Uses human readable YAML config files
- Has single node multi-GPU support
- MLFlow/TensorBoard integration
- Modular and extensible for adding custom image formats/layers/losses/metrics/callbacks /augmentations/preprocessing functionality
- Can import ImageryDataset class for use in your own framework to leverage DELTA's imagery handling anywhere you used Tensorflow Dataset

Example: Mapping clouds with Landsat-8 images

- Using USGS SPARCS dataset – Landsat 8
- Demonstrate:
  - Ease of ingesting imagery and labels for ML setup
  - Ease of applying trained models to imagery
  - Ease and customizability for training and inference both for ML scientists and less experienced domain scientists
- Jupyterlab demo

Landsat 8 Classification Example

This example will walk you through an example of using DELTA to train a simple example model. You can use what you learn here to use DELTA on your own datasets and with your own model architectures. In this example you will:

- Download a dataset of images and labels
- Train a simple model using example configuration files
- Examine results of the trained model
- Make some changes to the configuration files and train a modified model
- Examine the results of the modified model

Downloading and Extracting Dataset

The dataset includes satellite images along with classification labels for different types of land cover (water, cloud, snow, etc.).

In [1]:

```
!echo "Downloading dataset."
!curl -O https://landsat.usgs.gov/cloud-validation/sparcs/l8cloudmasks.zip
```

```
Downloading dataset.
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 1483M  100 1483M    0     0  2746k      0  0:09:13  0:09:13 --:--:-- 5555k15  0:00:18  0:08:57 1938k1  0:09:05
2236k2681k      0     0  0:09:26  0:00:26  0:09:00 2682k9 3187k3 3289k6k      0  0:08:37  0:01:04  0:07:33 1407k2836k
0   0:08:55  0:01:08  0:07:47 1279k0      0  2477k      0  0:10:13  0:01:42  0:08:31 3007k 0  0:09:19  0:02:15  0:07:
04 4545k 2732k      0  0:09:15  0:02:25  0:06:50 2926k1  0:02:59  0:06:22 1946k0:09:23  0:03:03  0:06:20 2334k 0  0
:09:34  0:03:17  0:06:17  599k  0   0:09:39  0:03:20  0:06:19 1247k   0   2583k      0  0:09:48  0:03:35  0:06:13 197
9k09:49  0:03:50  0:05:59 2954k      0  0:09:51  0:04:22  0:05:29 3312k 2587k      0  0:09:47  0:04:26  0:05:21 3763
k      0  0:09:46  0:04:37  0:05:09 2739k9:37  0:04:49  0:04:48 3964k5:17  0:04:21 2398k     0  0:09:40  0:05:26
0:04:14 2520k2619k      0  0:09:40  0:05:47  0:03:53 2867k 0   0:09:38  0:05:51  0:03:3
9 3367k      0  0:09:48  0:07:00  0:02:48 3117k 2836k2k      0  0:09:37  0:07:43  0:01:54 3750k01:49 2507k  0  261
4k      0  0:09:41  0:08:01  0:01:40 2271k 0:08:14  0:01:23 3636k      0  0:09:31  0:08:38  0:00:53 4221k
```

Here we're extracting the dataset and organizing the images into folders. We are:

- setting aside two images in a folder called "validate" to test our model later
- moving the satellite images into a folder called "train"
- moving the classification labels into a folder called "labels"

In [2]:

```
!echo "Extracting dataset."
!unzip -q l8cloudmasks.zip
!mkdir validate
!mv sending/LC82290562014157LGN00_24_data.tif sending/LC82210662014229LGN00_18_data.tif validate/
!mkdir train
!mv sending/*_data.tif train/
!mkdir labels
!mv sending/*_mask.png labels/
```

```
Extracting dataset.
mkdir: validate: File exists
mkdir: train: File exists
mkdir: labels: File exists
```

In [3]:

```
!ls
```

```
example_screenshots          l8_cloud_train_parameters.yaml
l8_cloud.sh                  l8cloudmasks.zip
l8_cloud_dataset.yaml        labels
l8_cloud_dataset_water.yaml  sending
l8_cloud_example.ipynb       train
l8_cloud_train_network.yaml  validate
```

In [4]:

```
!ls train/ | head
```

```
LC80010812013365LGN00_18_data.tif
LC80020622013244LGN00_32_data.tif
LC80050152014172LGN00_12_data.tif
LC80050562014076LGN00_33_data.tif
LC80150242014146LGN00_23_data.tif
LC80190352014078LGN01_26_data.tif
LC80200462014213LGN00_11_data.tif
LC80250402013245LGN00_45_data.tif
LC80250482014072LGN00_18_data.tif
LC80310432013207LGN00_11_data.tif
```

Example Satellite Image

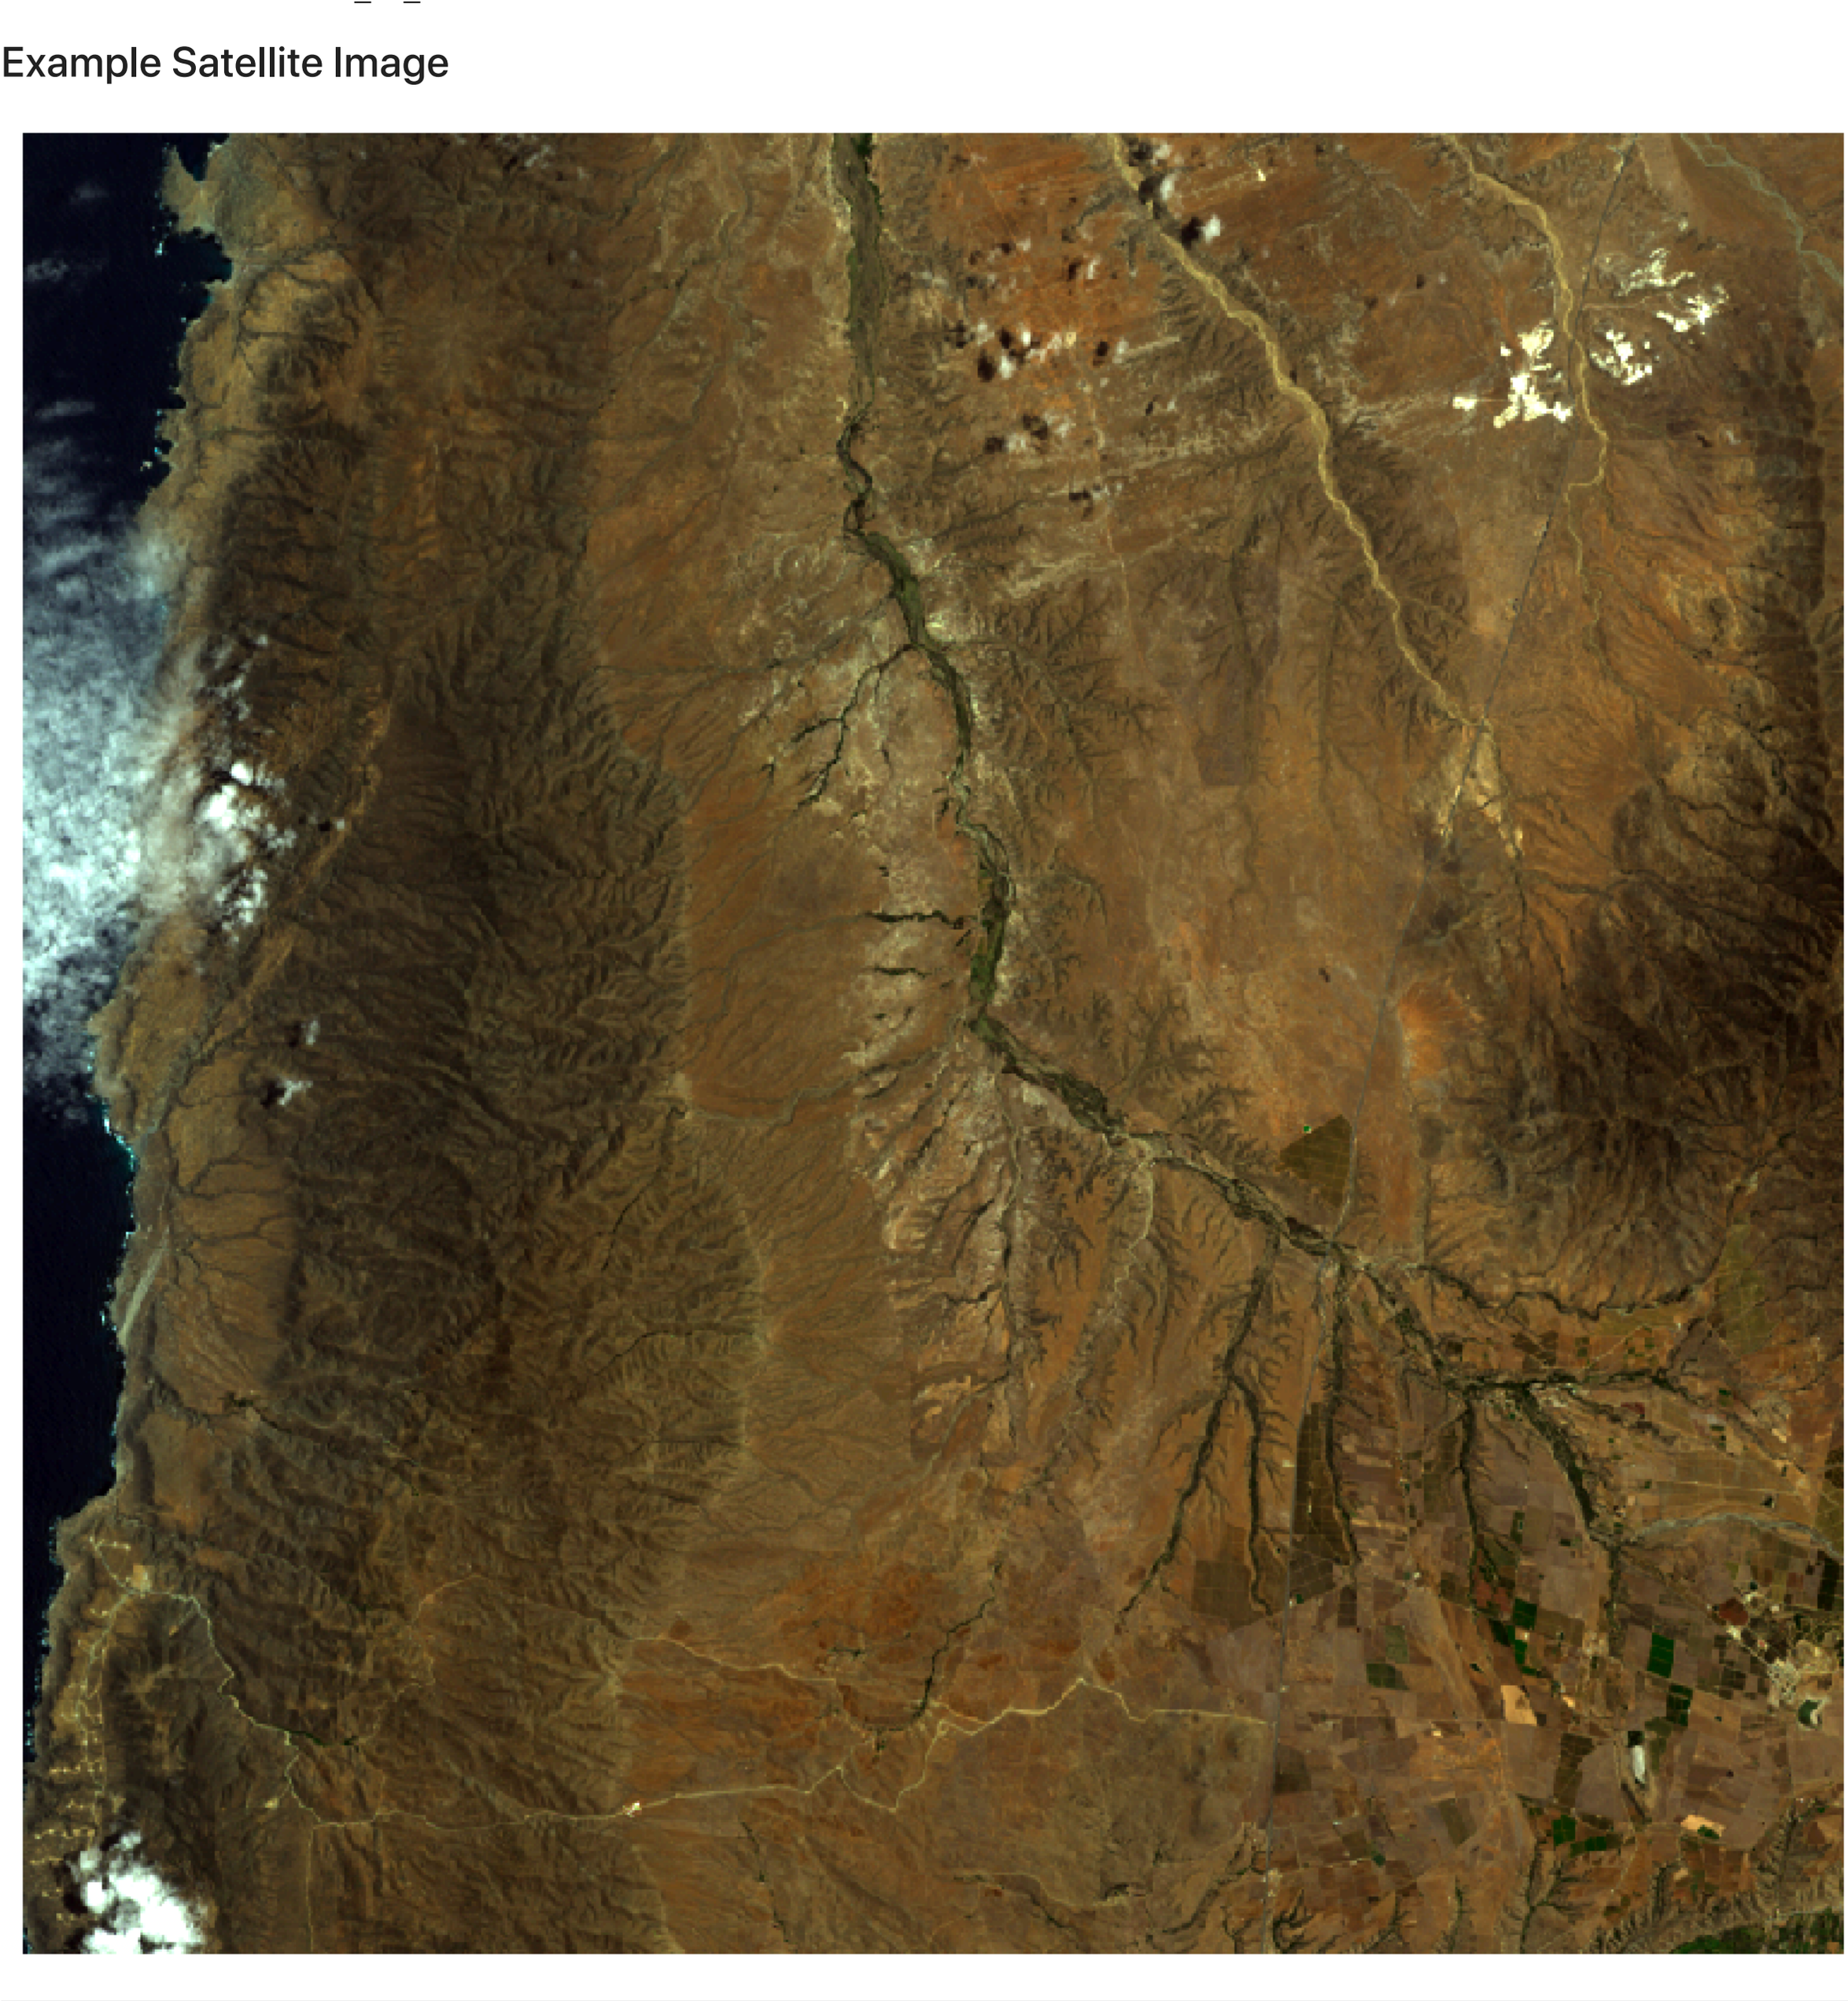

In [5]:

```
!ls labels | head
```

```
LC80010812013365LGN00_18_mask.png
LC80020622013244LGN00_32_mask.png
LC80050152014172LGN00_12_mask.png
LC80050562014076LGN00_33_mask.png
LC80150242014146LGN00_23_mask.png
LC80190352014078LGN01_26_mask.png
LC80200462014213LGN00_11_mask.png
LC80250402013245LGN00_45_mask.png
LC80250482014072LGN00_18_mask.png
LC80310432013207LGN00_11_mask.png
```

Example Label Image

The different colors represent different land cover classifications.

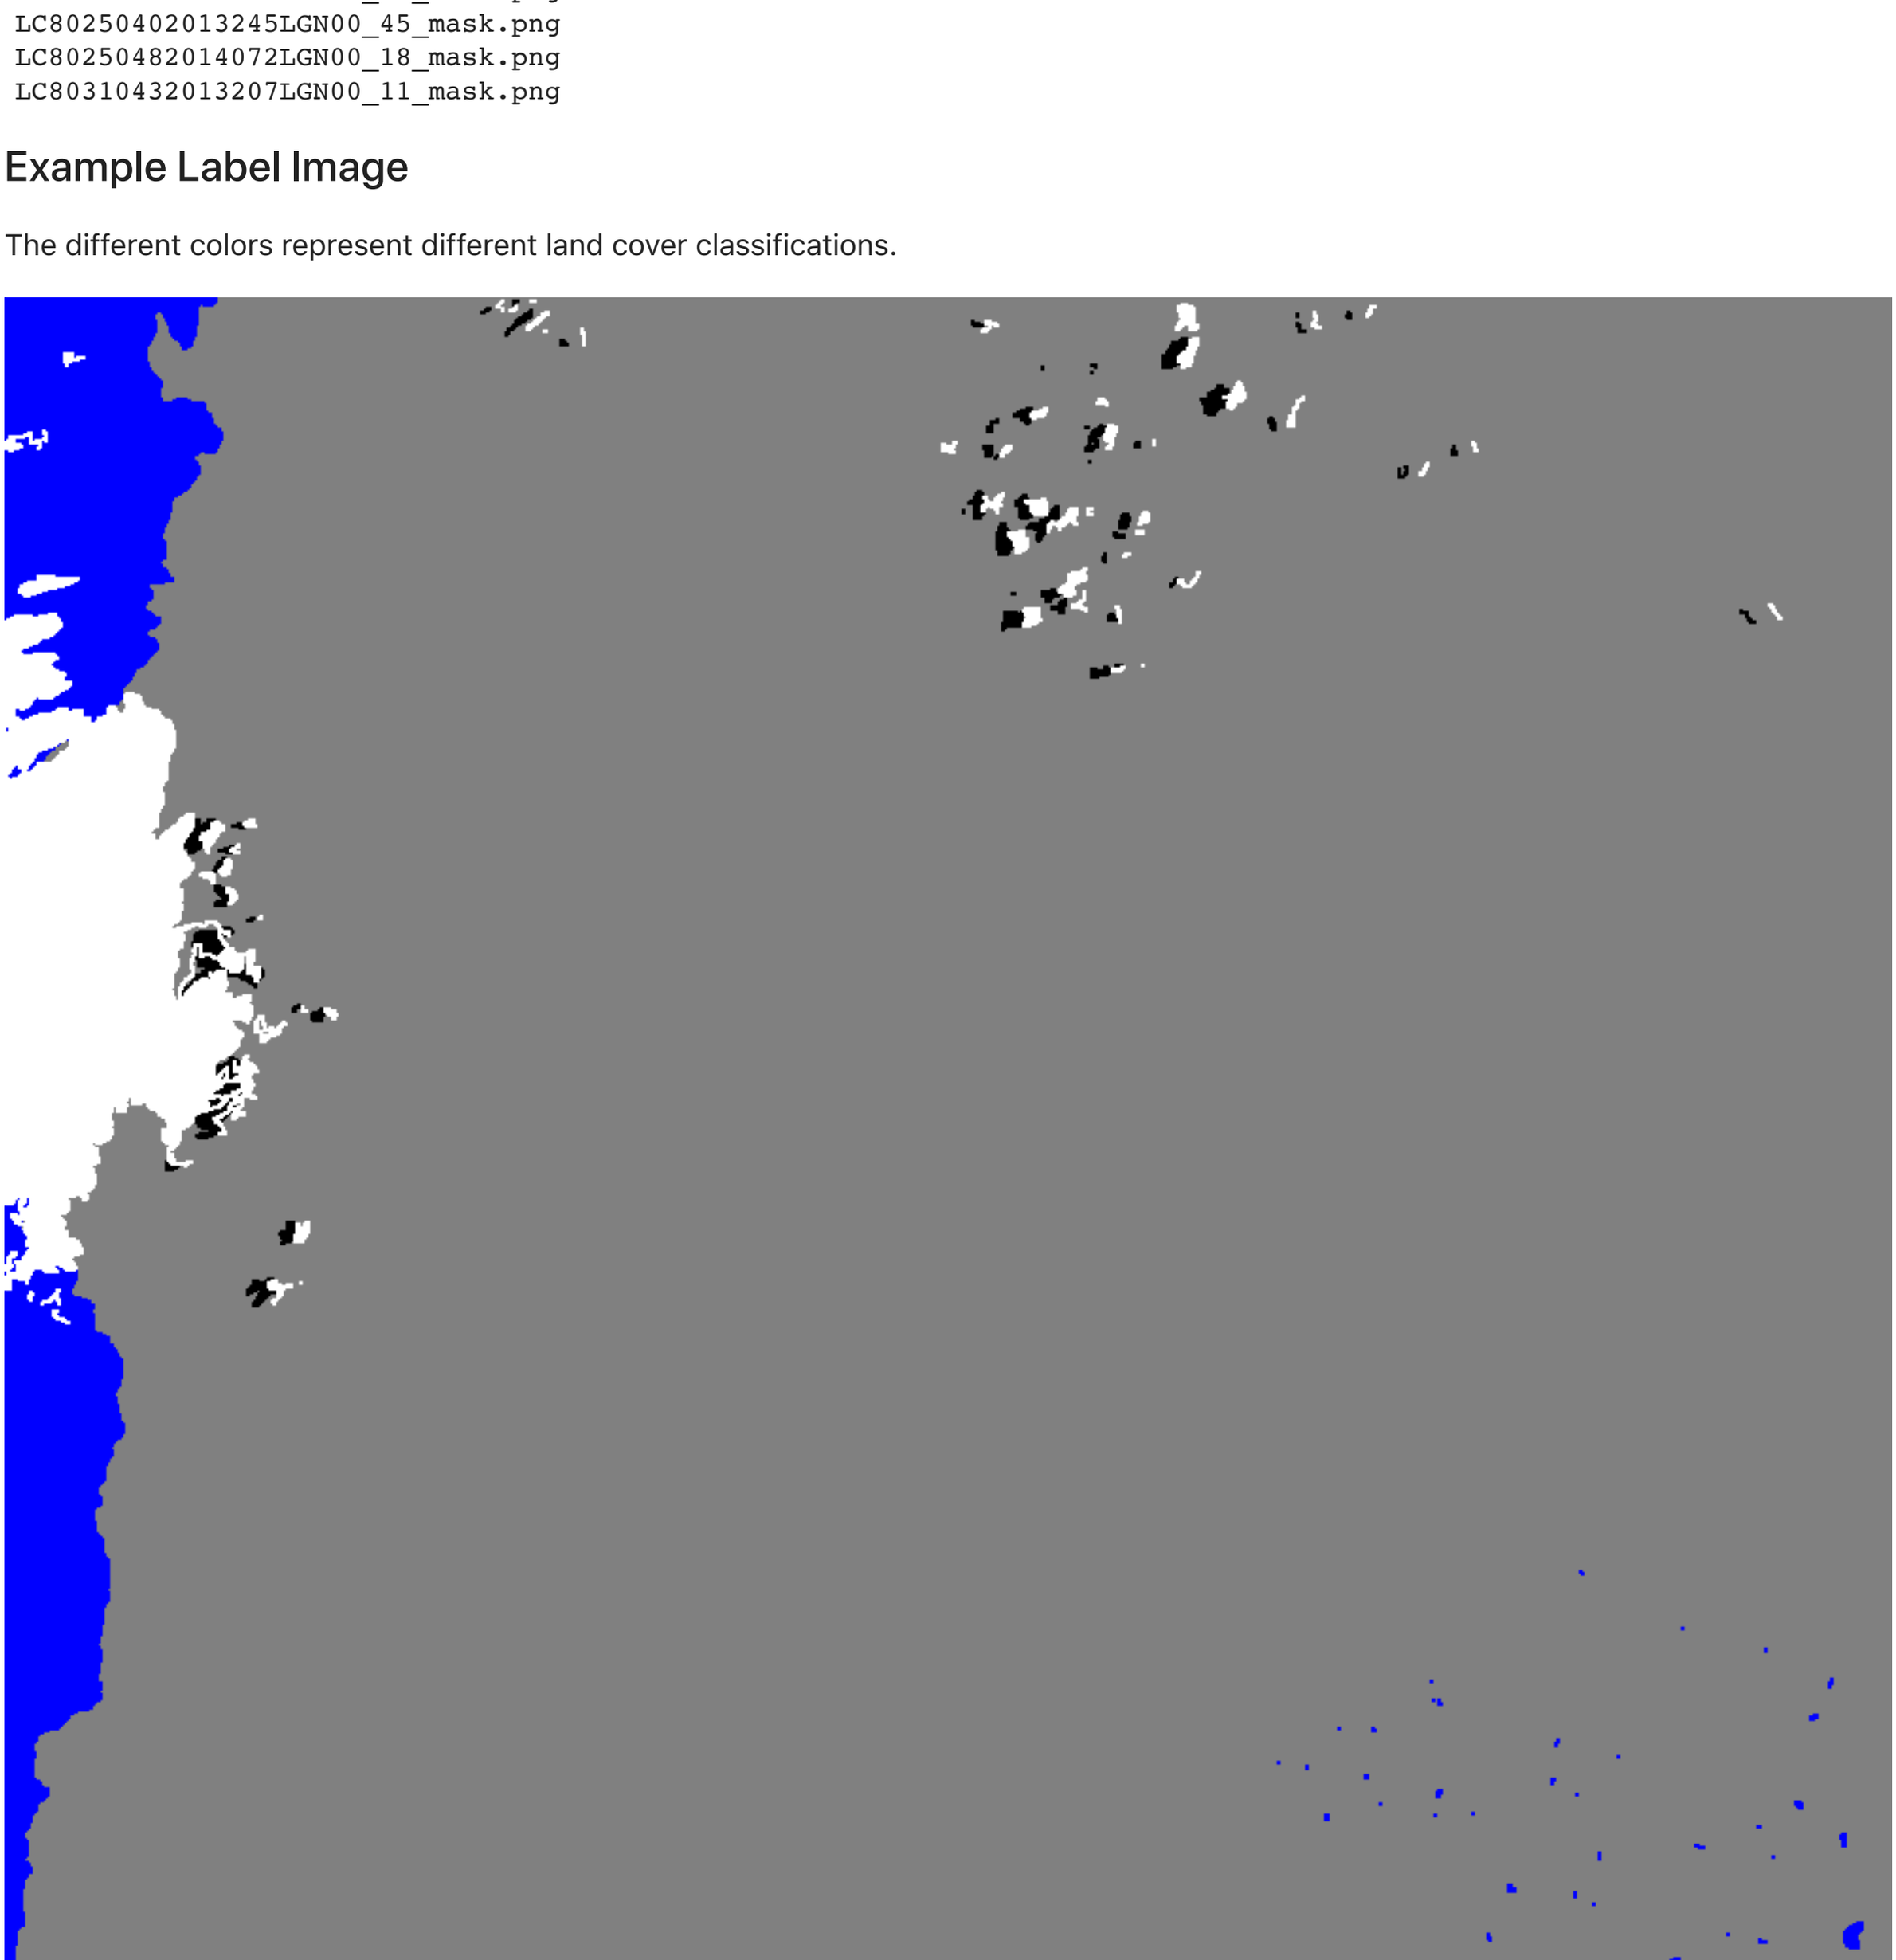

Training a Model

We'll use the following configuration files which specify a dataset to use (l8_cloud_dataset.yaml), a network architecture to train (l8_cloud_train_network.yaml), and the parameters to train with (l8_cloud_train_parameters.yaml).

Configuration YAMLs

l8_cloud_dataset.yaml
l8_cloud_train_network.yaml
l8_cloud_train_parameters.yaml

Detailed Config Documentation

In [6]:

```
!delta train --config l8_cloud_dataset.yaml --config l8_cloud_train_network.yaml --config l8_cloud_train_parameters
```

```
Epoch 1/10
31/31 [==============================] - 304s 10s/step - loss: 1.1059 - sparse_categorical_accuracy: 0.6334
Epoch 2/10
31/31 [==============================] - 248s 8s/step - loss: 0.6729 - sparse_categorical_accuracy: 0.7778
Epoch 3/10
31/31 [==============================] - 244s 8s/step - loss: 0.6341 - sparse_categorical_accuracy: 0.7914
Epoch 4/10
31/31 [==============================] - 254s 8s/step - loss: 0.6193 - sparse_categorical_accuracy: 0.8045
Epoch 5/10
31/31 [==============================] - 273s 9s/step - loss: 0.5835 - sparse_categorical_accuracy: 0.8108
Epoch 6/10
31/31 [==============================] - 250s 8s/step - loss: 0.5541 - sparse_categorical_accuracy: 0.8186
Epoch 7/10
31/31 [==============================] - 240s 8s/step - loss: 0.5554 - sparse_categorical_accuracy: 0.8185
Epoch 8/10
31/31 [==============================] - 233s 8s/step - loss: 0.5567 - sparse_categorical_accuracy: 0.8190
Epoch 9/10
31/31 [==============================] - 250s 8s/step - loss: 0.5207 - sparse_categorical_accuracy: 0.8308
Epoch 10/10
31/31 [==============================] - 246s 8s/step - loss: 0.5108 - sparse_categorical_accuracy: 0.8438

Finished, saving model to file:///Users/mvonpohl/Library/Application Support/delta/mlflow/1/db96fb5d0418440786a7b0b
ee1696732/artifacts/final_model.savedmodel.
Elapsed time =  2650.9034371376038
```

Examine Model Results

The previous step produced a trained model. Now we can use the model to classify the images we set aside in the "validate" folder.

In [12]:

```
!delta classify --config l8_cloud_dataset.yaml --image-dir ./validate --outdir ./model_output --overlap 32 l8_clou
```

```
LC82290562014157LGN00_24_data : |██████████████████████████████|
9 / 9
Image: ./validate/LC82290562014157LGN00_24_data.tif
Shadow               --- Precision:  58.43%    Recall:  39.25%        Frequency:  13.64%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   3.84%
Water                --- Precision:  77.45%    Recall:  94.50%        Frequency:  24.57%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:  75.19%    Recall:  96.38%        Frequency:  29.91%
Cloud                --- Precision:  96.24%    Recall:  77.30%        Frequency:  28.05%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 79.07% Accuracy

LC82210662014229LGN00_18_data : |██████████████████████████████|
9 / 9
Image: ./validate/LC82210662014229LGN00_18_data.tif
Shadow               --- Precision:  88.75%    Recall:  20.58%        Frequency:  10.30%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.02%
Water                --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:  87.34%    Recall:  99.36%        Frequency:  75.67%
Cloud                --- Precision:  94.27%    Recall:  63.73%        Frequency:  14.02%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 86.24% Accuracy

Overall:
Shadow               --- Precision:  64.70%    Recall:  31.22%        Frequency:  11.97%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   1.93%
Water                --- Precision:  72.48%    Recall:  94.50%        Frequency:  12.28%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:  83.60%    Recall:  98.52%        Frequency:  52.79%
Cloud                --- Precision:  95.65%    Recall:  72.78%        Frequency:  21.03%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 82.65% Accuracy

Elapsed time =  5.5682148933410645
```

Validation Image

The validation image we fed into the model.

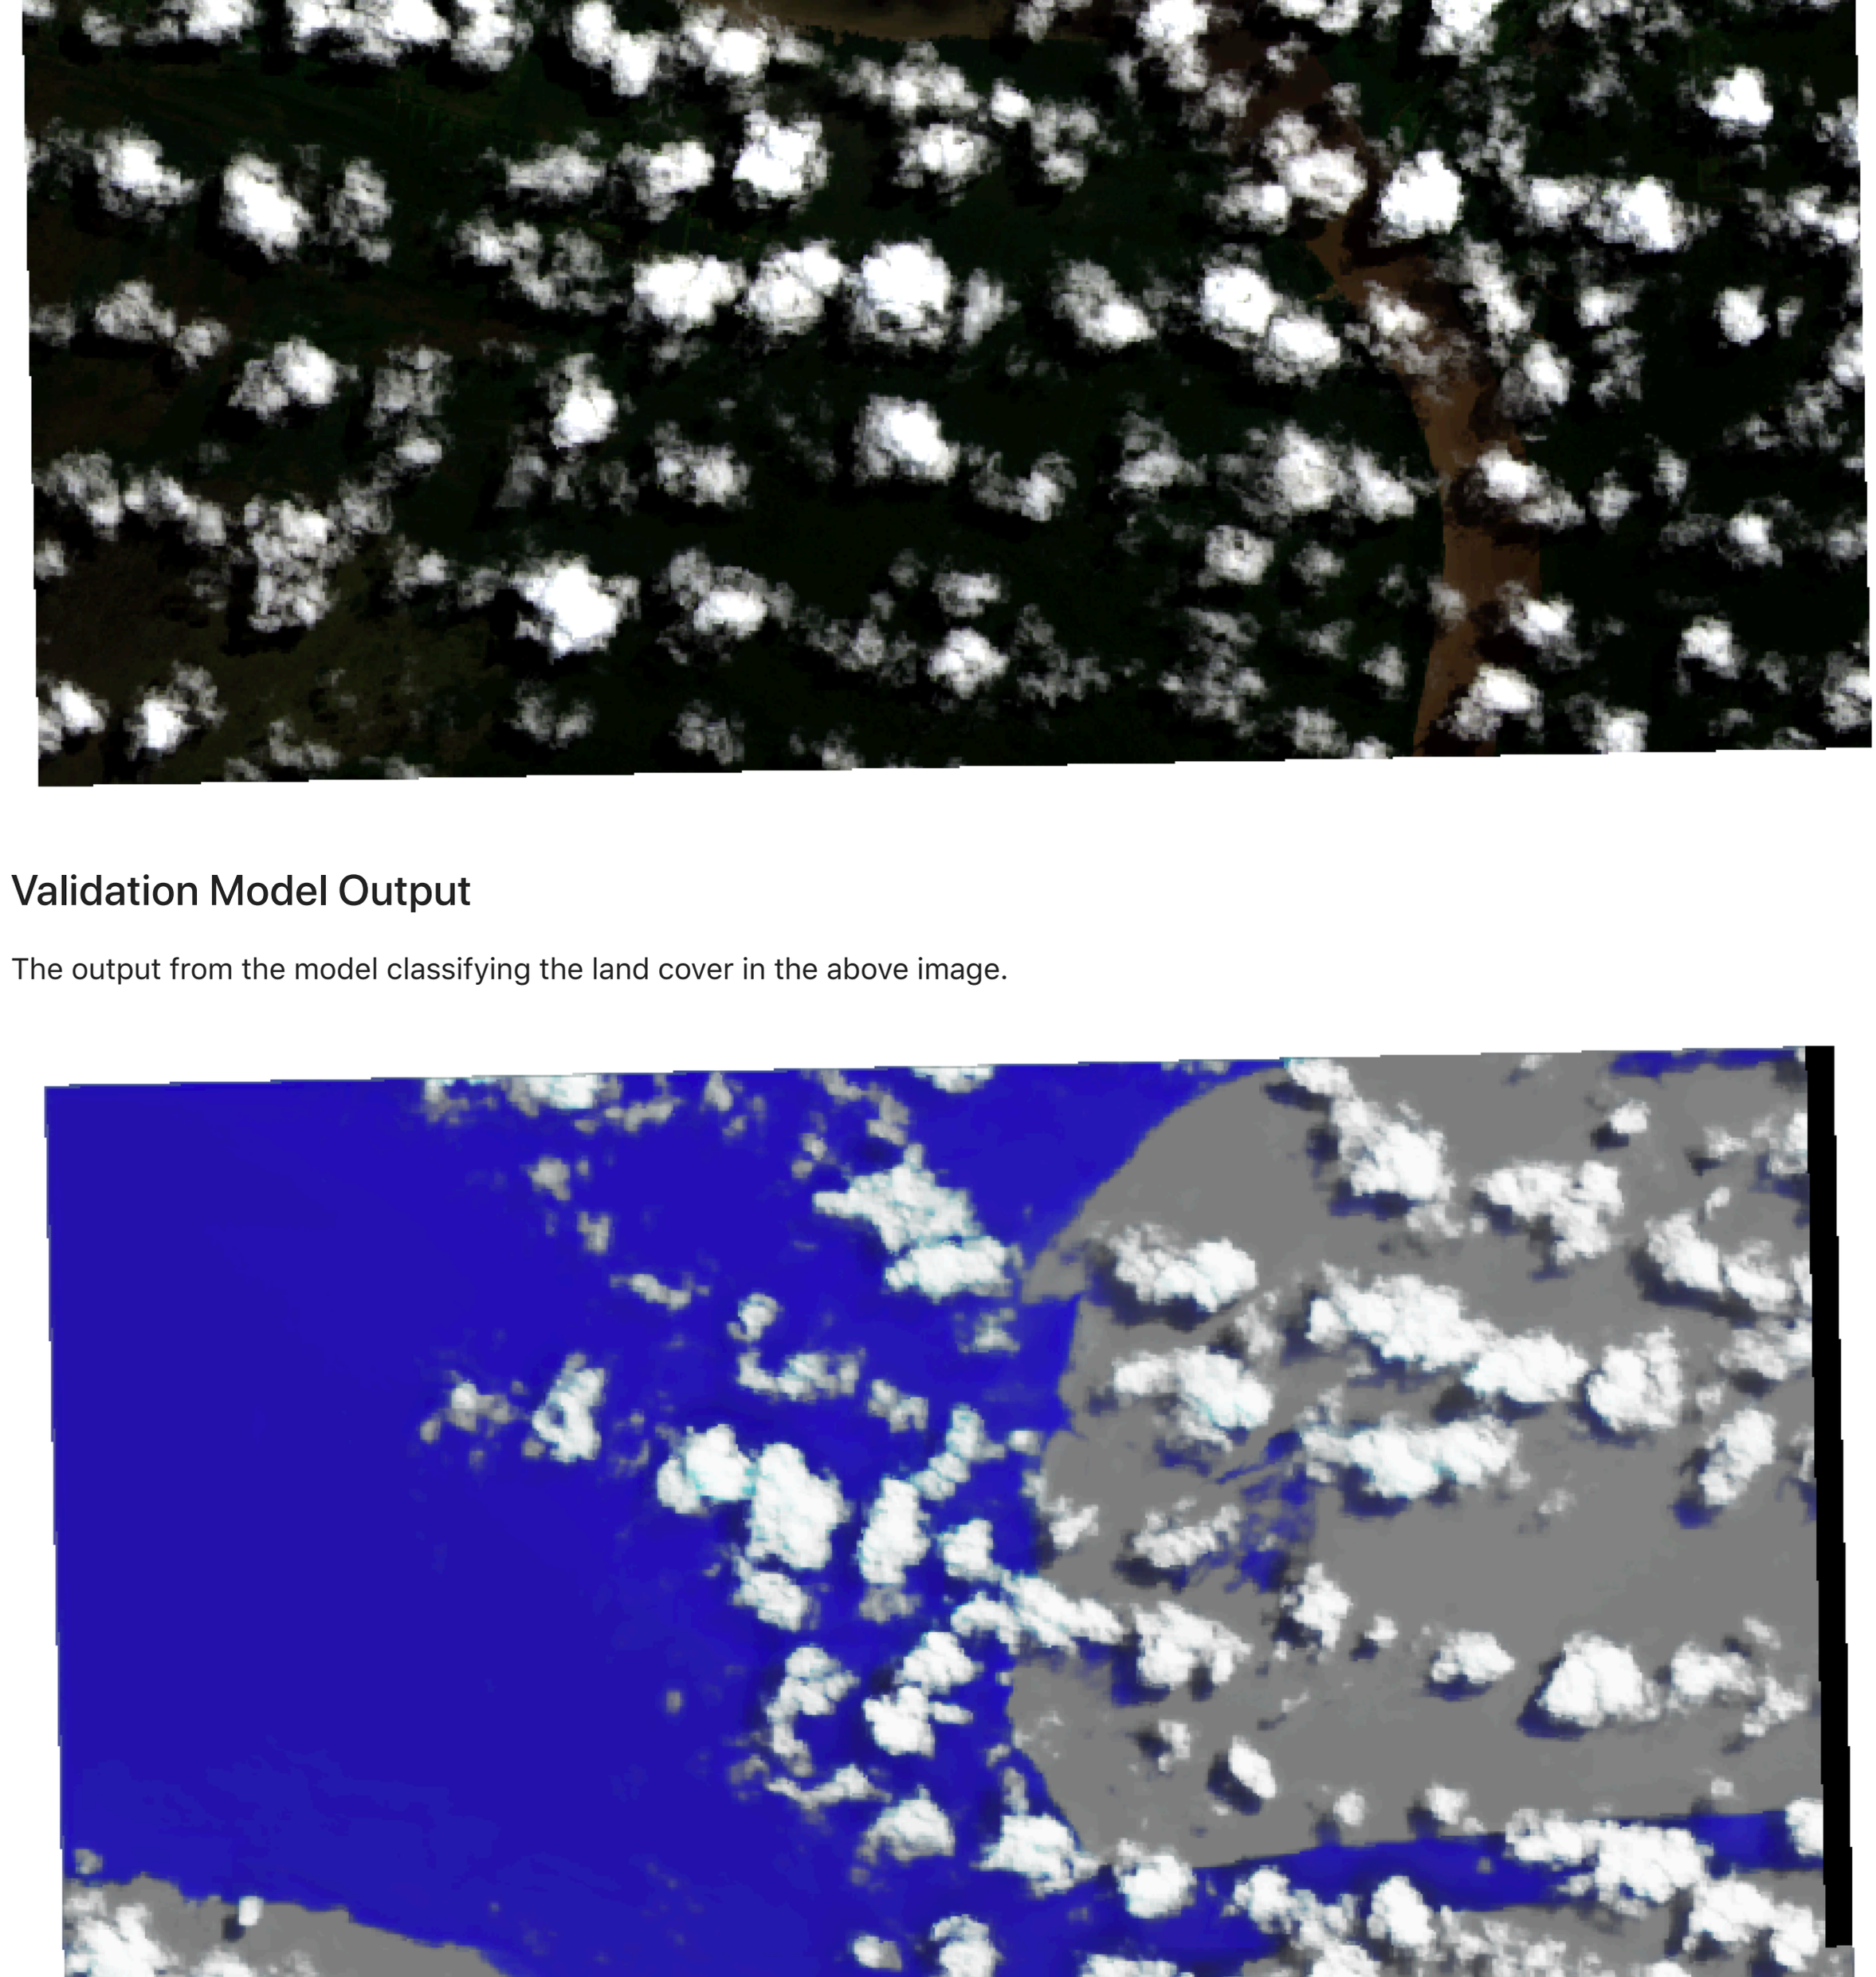

Validation Model Output

The output from the model classifying the land cover in the above image.

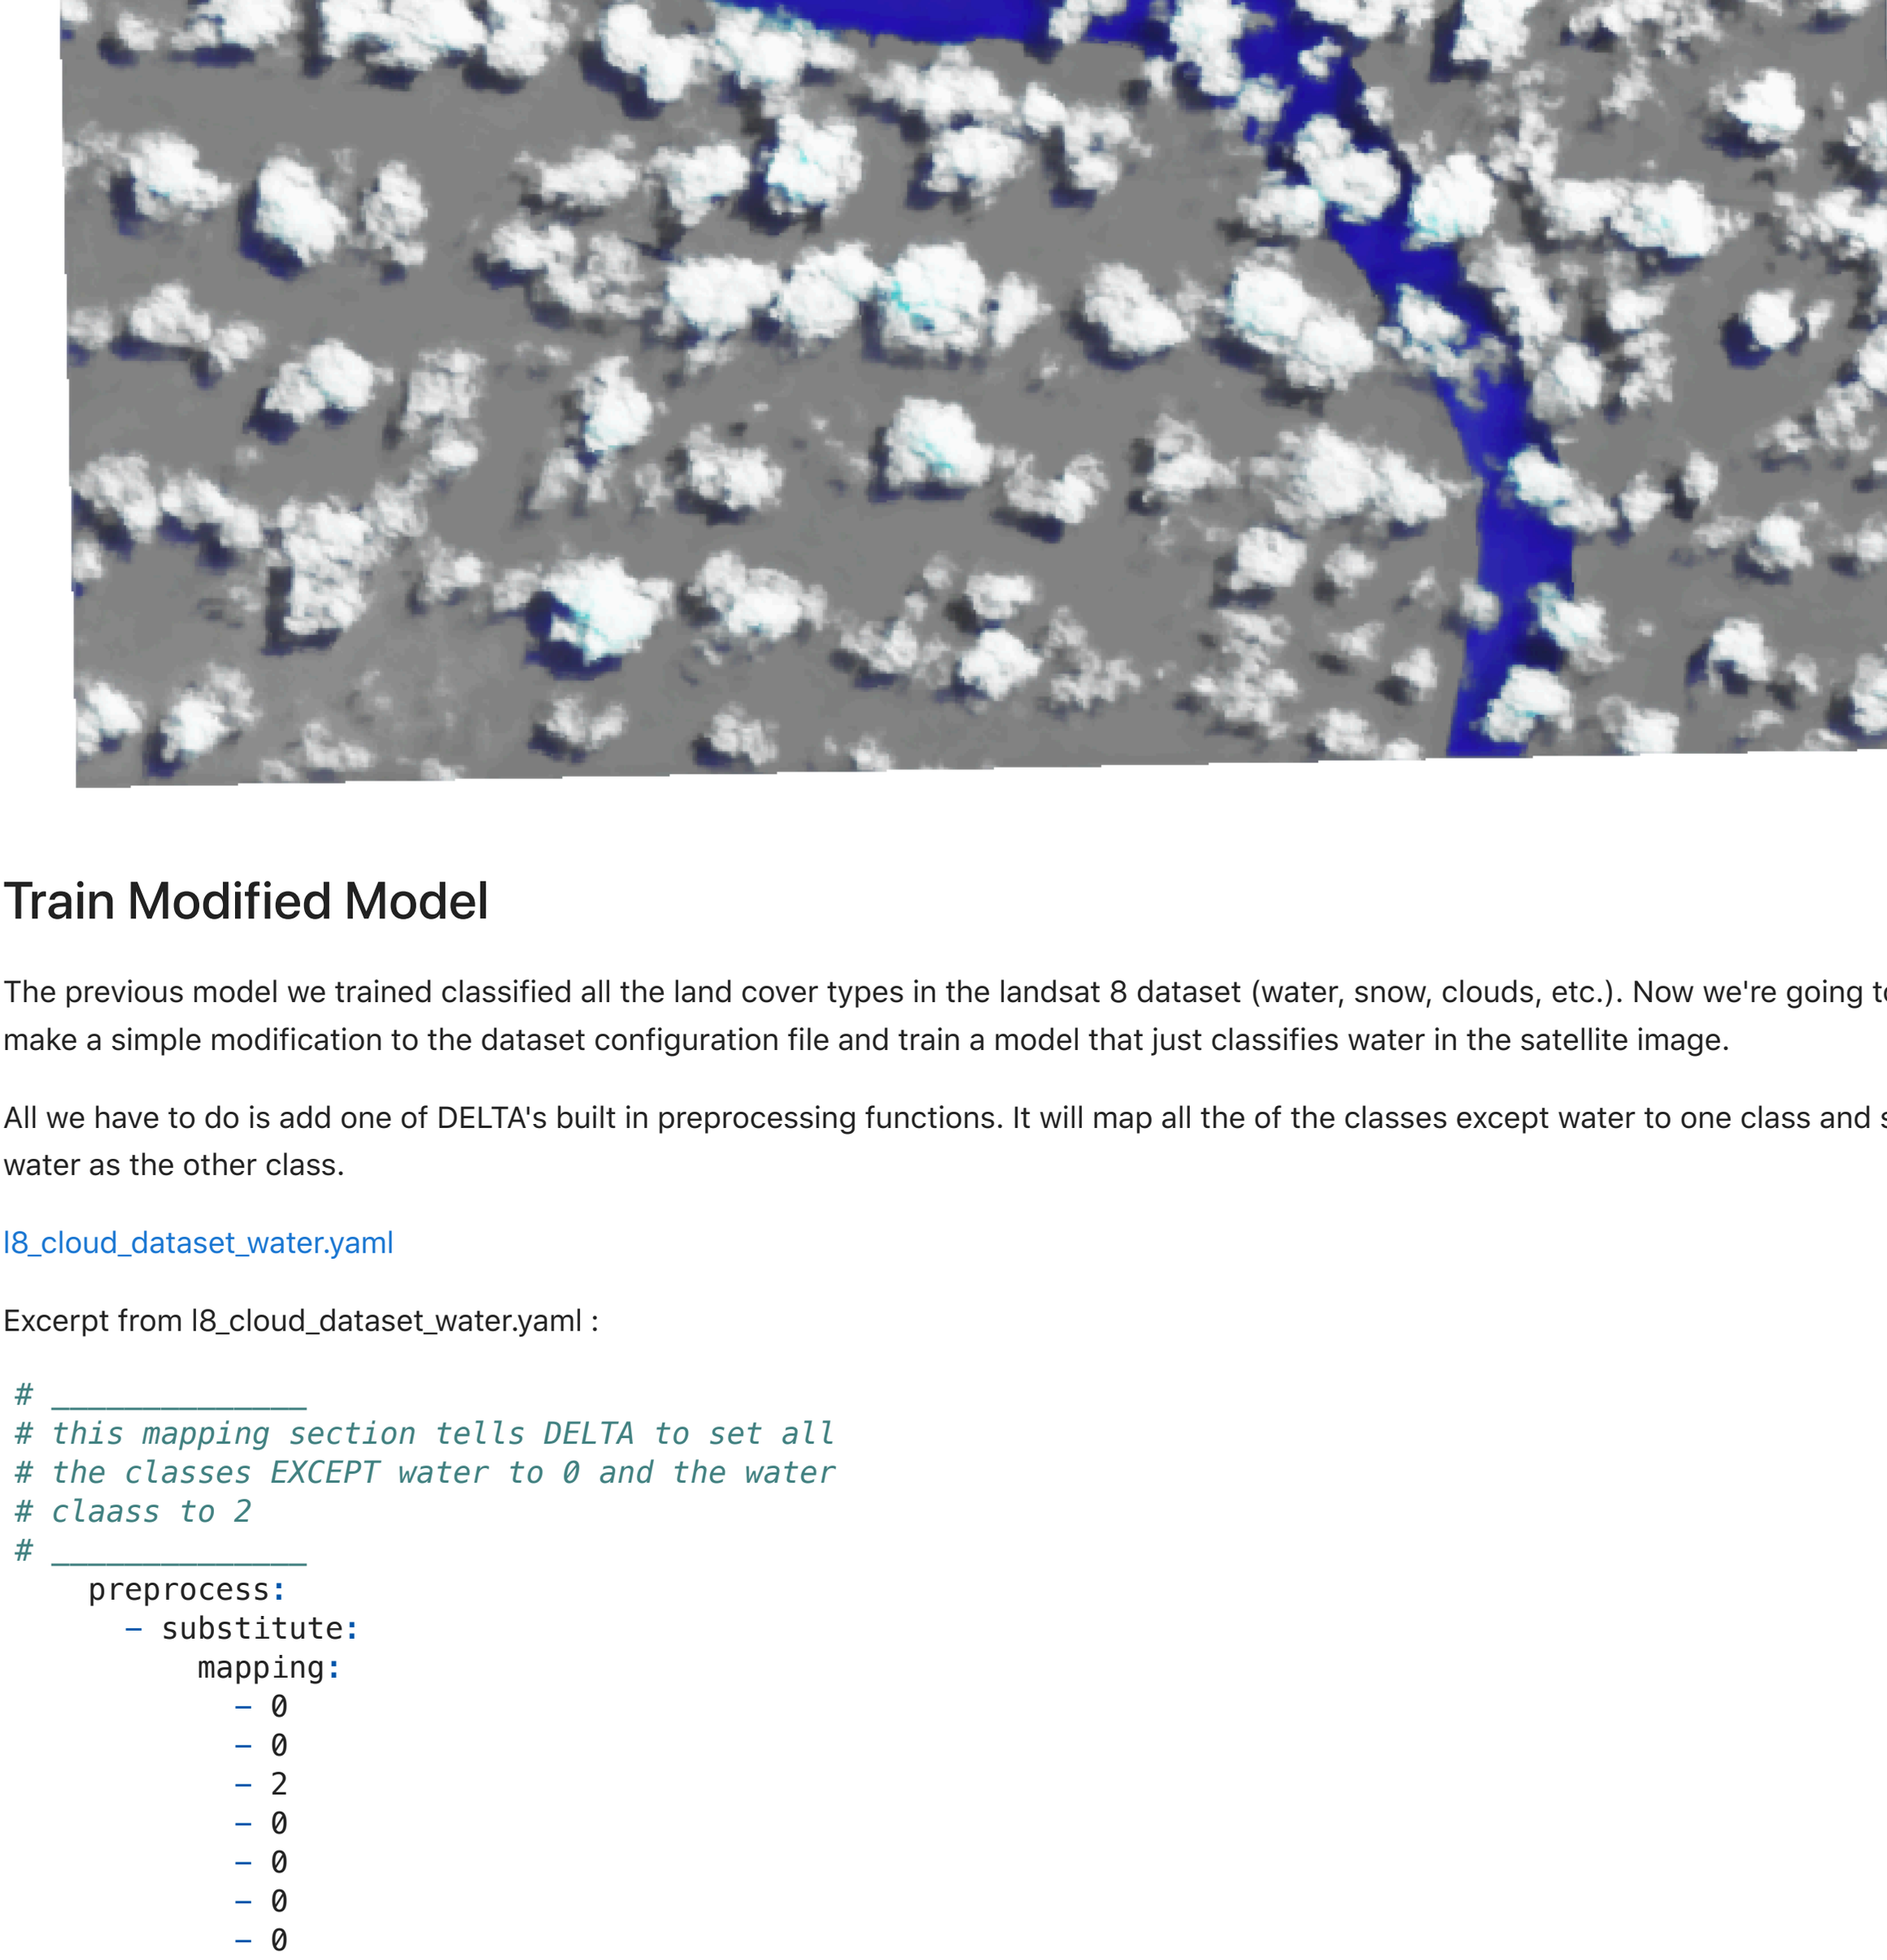

Train Modified Model

The previous model we trained classified all the land cover types in the landsat 8 dataset (water, snow, clouds, etc.). Now we're going to make a simple modification to the dataset configuration file and train a model that just classifies water in the satellite image.

All we have to do is add one of DELTA's built in preprocessing functions. It will map all the of the classes except water to one class and set water as the other class.

l8_cloud_dataset_water.yaml

Excerpt from l8_cloud_dataset_water.yaml :

```
# ______________
# this mapping section tells DELTA to set all
# the classes EXCEPT water to 0 and the water
# claass to 2
# ______________
    preprocess:
      - substitute:
          mapping:
            - 0
            - 0
            - 2
            - 0
            - 0
            - 0
            - 0
  classes:
    - 0:
        name: Shadow
        color: 0x000000
    - 1:
        name: Shadow over Water
        color: 0x000080
    - 2:
        name: Water
        color: 0x0000FF
    - 3:
        name: Snow
        color: 0x00FFFF
    - 4:
        name: Land
        color: 0x808080
    - 5:
        name: Cloud
        color: 0xFFFFFF
    - 6:
        name: Flooded
        color: 0x808000
```

In [8]:

```
!delta train --config l8_cloud_dataset_water.yaml --config l8_cloud_train_network.yaml --config l8_cloud_train_para
```

```
Epoch 1/10
31/31 [==============================] - 227s 7s/step - loss: 0.7085 - sparse_categorical_accuracy: 0.8178
Epoch 2/10
31/31 [==============================] - 236s 8s/step - loss: 0.1522 - sparse_categorical_accuracy: 0.9524
Epoch 3/10
31/31 [==============================] - 233s 8s/step - loss: 0.1161 - sparse_categorical_accuracy: 0.9603
Epoch 4/10
31/31 [==============================] - 238s 8s/step - loss: 0.1055 - sparse_categorical_accuracy: 0.9650
Epoch 5/10
31/31 [==============================] - 231s 7s/step - loss: 0.1007 - sparse_categorical_accuracy: 0.9688
Epoch 6/10
31/31 [==============================] - 236s 8s/step - loss: 0.0889 - sparse_categorical_accuracy: 0.9701
Epoch 7/10
31/31 [==============================] - 228s 7s/step - loss: 0.1051 - sparse_categorical_accuracy: 0.9635
Epoch 8/10
31/31 [==============================] - 231s 7s/step - loss: 0.1050 - sparse_categorical_accuracy: 0.9602
Epoch 9/10
31/31 [==============================] - 230s 7s/step - loss: 0.0893 - sparse_categorical_accuracy: 0.9695
Epoch 10/10
31/31 [==============================] - 247s 8s/step - loss: 0.0963 - sparse_categorical_accuracy: 0.9635

Finished, saving model to file:///Users/mvonpohl/Library/Application Support/delta/mlflow/1/f06e99d8150f42a18eab169
67ee94f95/artifacts/final_model.savedmodel.
Elapsed time =  2443.023379802704
```

Examine Model Results

Now we can examine the new model by classifying the images we set aside in the "validate" folder.

In [13]:

```
!delta classify --config l8_cloud_dataset_water.yaml --image-dir ./validate --outdir ./model_output_water --overlap
```

```
LC82290562014157LGN00_24_data : |██████████████████████████████|
9 / 9
Image: ./validate/LC82290562014157LGN00_24_data.tif
Shadow               --- Precision:  98.83%    Recall:  93.78%        Frequency:  75.43%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Water                --- Precision:  83.48%    Recall:  96.59%        Frequency:  24.57%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Cloud                --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 94.47% Accuracy

LC82210662014229LGN00_18_data : |██████████████████████████████|
9 / 9
Image: ./validate/LC82210662014229LGN00_18_data.tif
Shadow               --- Precision: 100.00%    Recall:  99.92%        Frequency: 100.00%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Water                --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Cloud                --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 99.92% Accuracy

Overall:
Shadow               --- Precision:  99.51%    Recall:  97.28%        Frequency:  87.72%
Shadow over Water    --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Water                --- Precision:  83.25%    Recall:  96.59%        Frequency:  12.28%
Snow                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Land                 --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Cloud                --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
Flooded              --- Precision:   0.00%    Recall:   0.00%        Frequency:   0.00%
 97.19% Accuracy

Elapsed time =  4.435237169265747
```

Validation Image

The validation image we fed into the model.

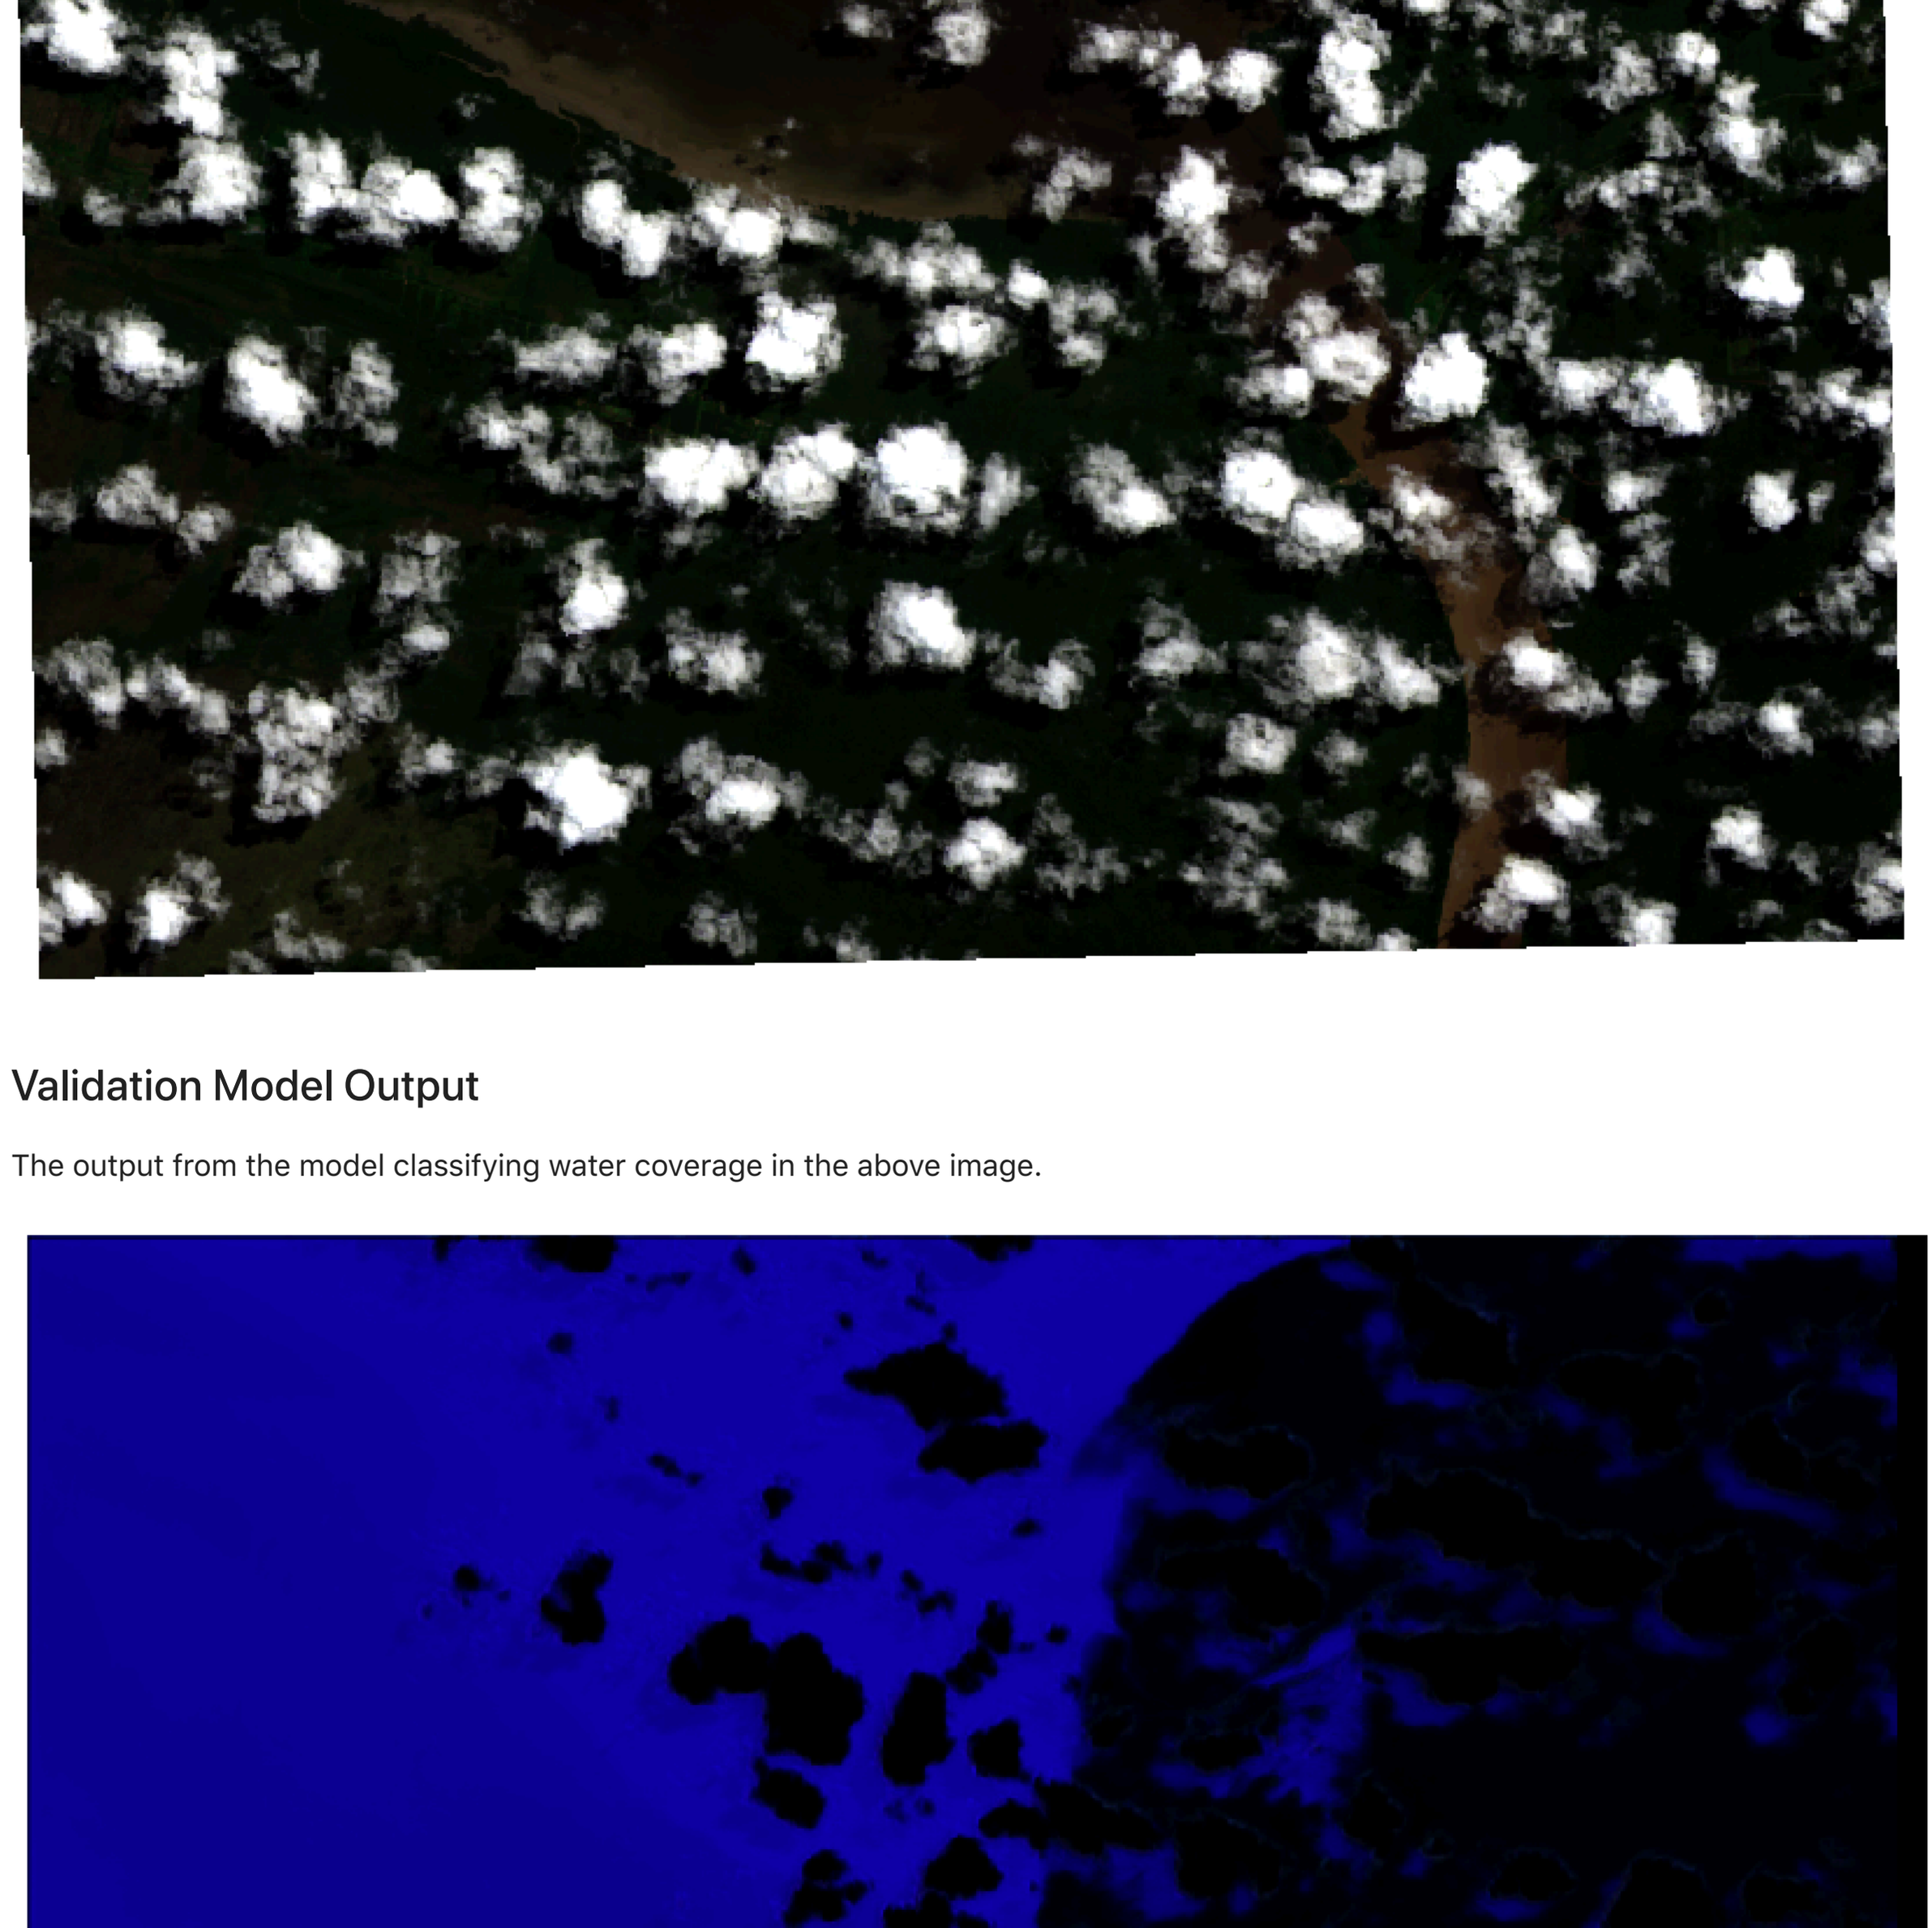

Validation Model Output

The output from the model classifying water coverage in the above image.

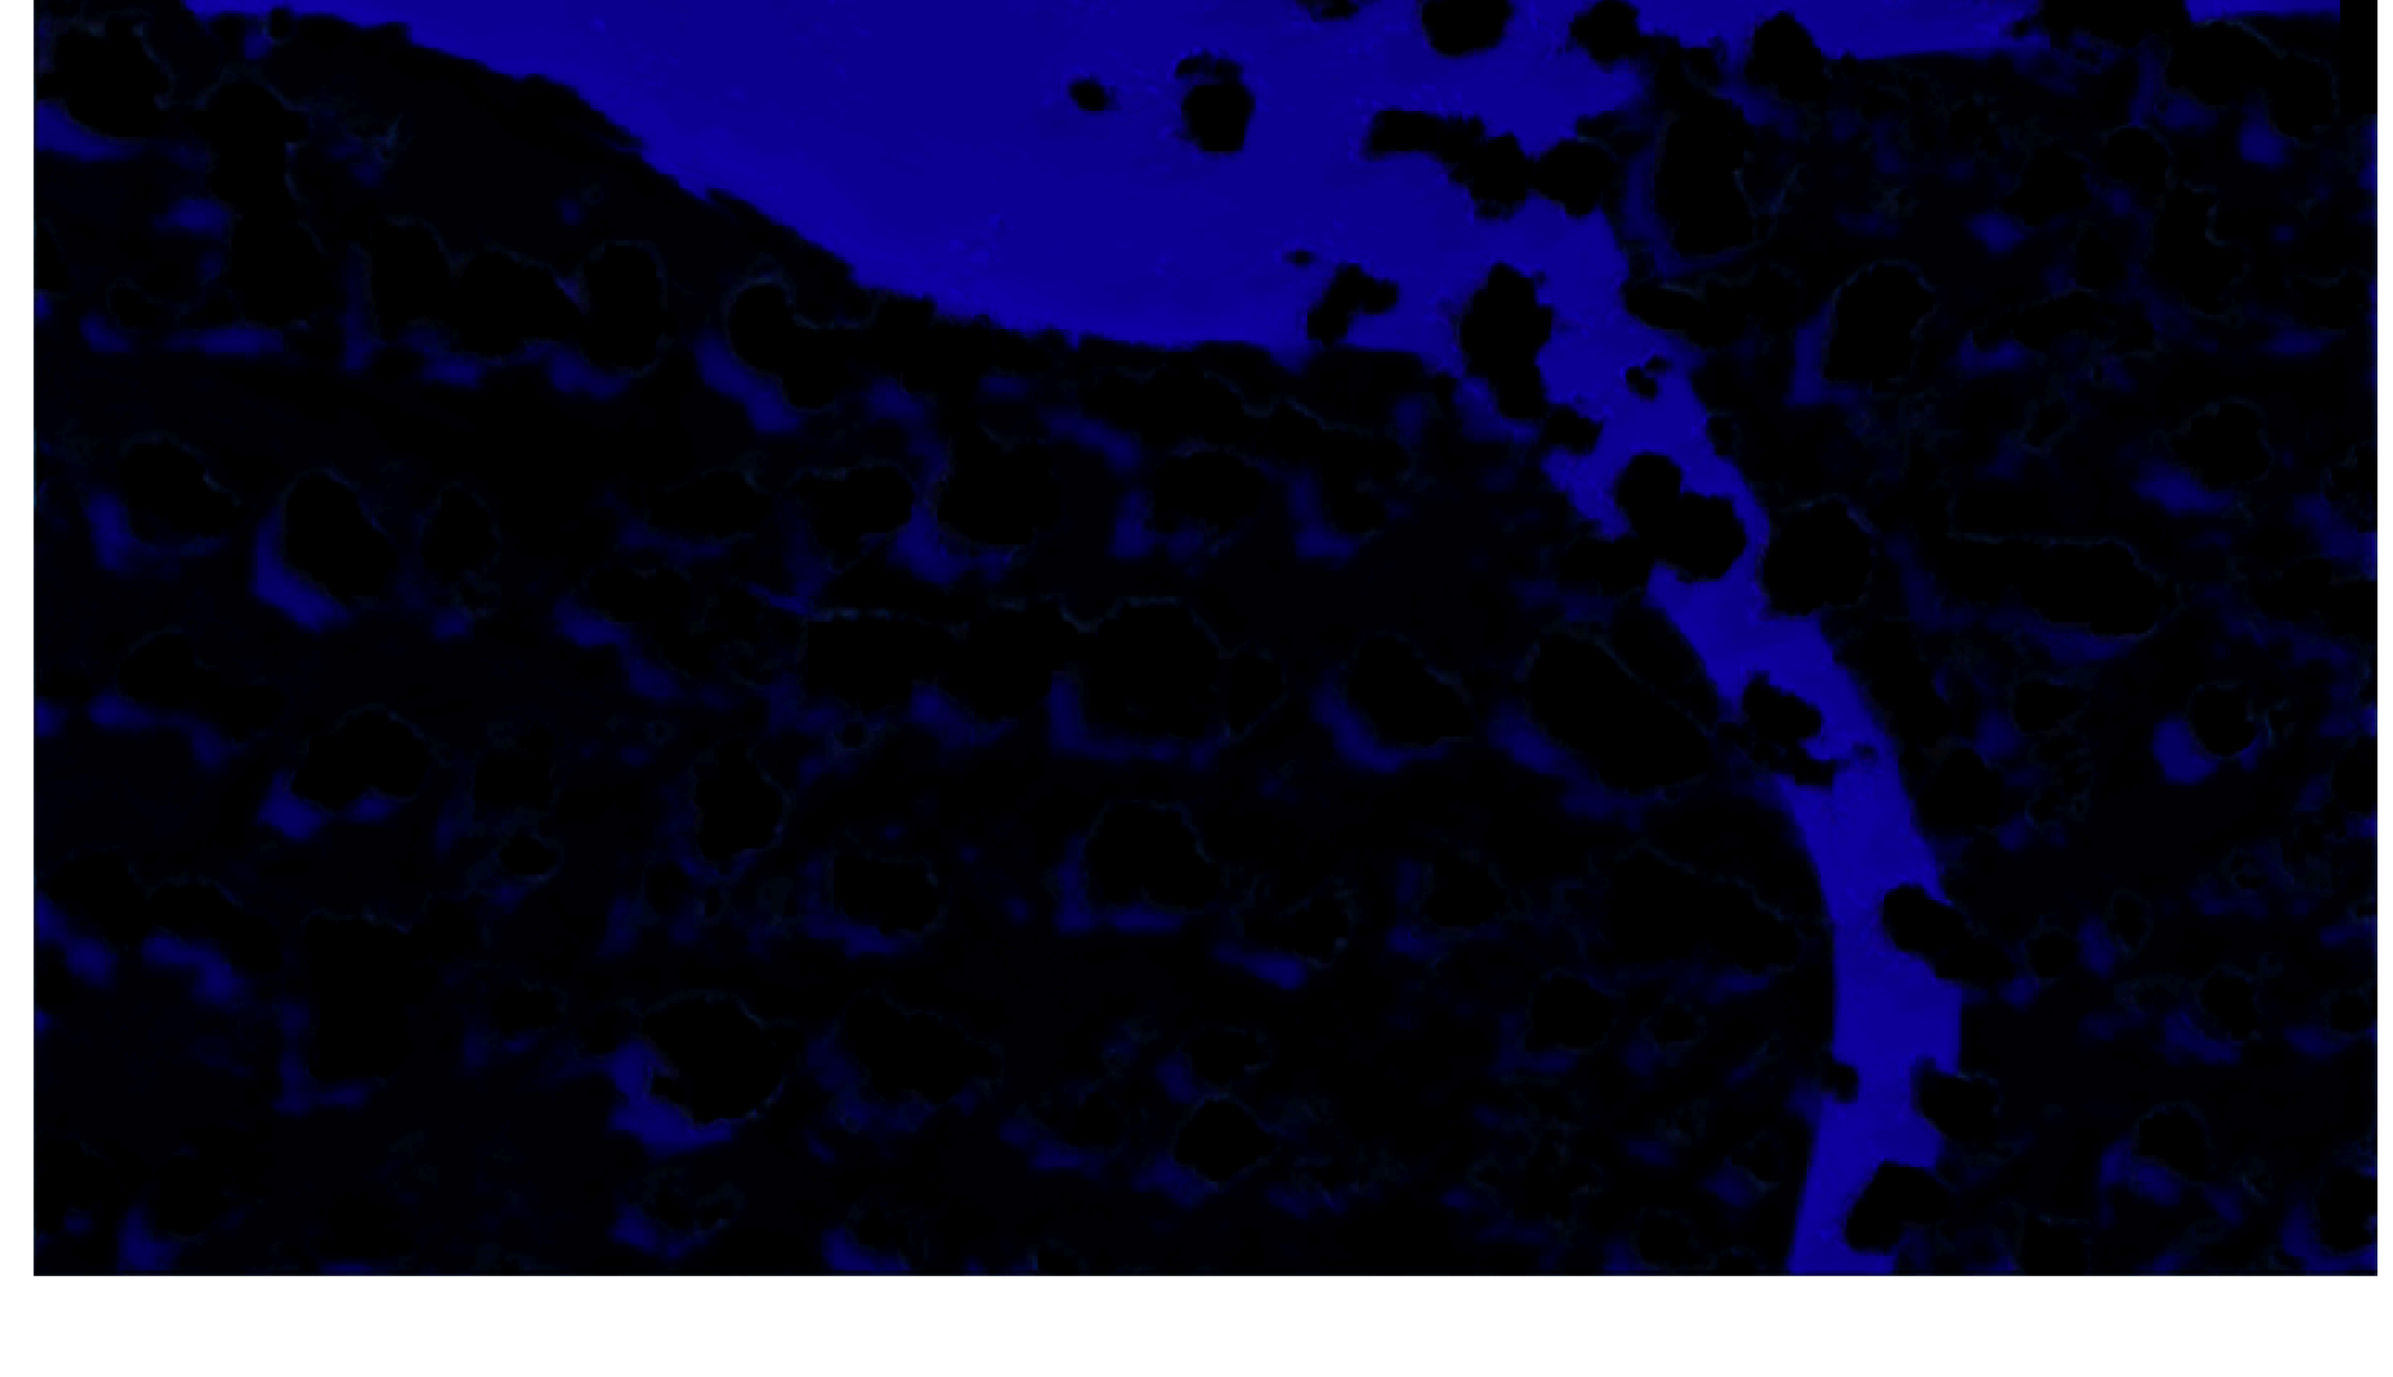

DELTA

- https://github.com/nasa/delta

- Michael von Pohle – michael.vonpohle@nasa.gov

```
train:
  network:
    layers:
     - Input:
         shape: [~, ~, num_bands]
     - Conv2D:
          filters: 16
          kernel_size: [3, 3]
          padding: same
     - BatchNormalization:
     - Activation:
         activation: relu
         name: c1
     - Dropout:
         rate: 0.2
     - MaxPool2D:
     - Conv2D:
          filters: 32
          kernel_size: [3, 3]
          padding: same
     - BatchNormalization:
     - Activation:
         activation: relu
         name: c2
     - Dropout:
         rate: 0.2
     - MaxPool2D:
     - Conv2D:
          filters: 64
          kernel_size: [3, 3]
          padding: same
     - BatchNormalization:
     - Activation:
         activation: relu
         name: c3
     - Dropout:
         rate: 0.2
     - MaxPool2D:
     - Conv2D:
          filters: 128
          kernel_size: [3, 3]
          padding: same
     - BatchNormalization:
     - Activation:
         activation: relu
         name: c4
     - UpSampling2D:
     - Conv2D:
          filters: 64
          kernel_size: [2, 2]
          padding: same
     - BatchNormalization:
     - Activation:
         activation: relu
         name: u3
     - Concatenate:
         inputs: [c3, u3]
     - Dropout:
         rate: 0.2
     - Conv2D:
          filters: 64
```

```
      kernel_size: [3, 3]
      padding: same
- UpSampling2D:
- Conv2D:
      filters: 32
      kernel_size: [2, 2]
      padding: same
- BatchNormalization:
- Activation:
    activation: relu
    name: u2
- Concatenate:
    inputs: [c2, u2]
- Dropout:
    rate: 0.2
- Conv2D:
      filters: 32
      kernel_size: [3, 3]
      padding: same
- UpSampling2D:
- Conv2D:
      filters: 16
      kernel_size: [2, 2]
      padding: same
- BatchNormalization:
- Activation:
    activation: relu
    name: u1
- Concatenate:
    inputs: [c1, u1]
- Dropout:
    rate: 0.2
- Conv2D:
    filters: 7
    kernel_size: [3, 3]
    activation: linear
    padding: same
- Softmax:
    axis: 3
```