SolarX:
Solar Panel Segmentation and Classification

Authors: Spencer Paul, Ethan Hellman, Rodri Neito

Background

Solar photovoltaics (PV) is an exponentially growing form of renewable energy and many countries have been making efforts to install solar cells on rooftops. However, the scaling of PVs is limited by land availability and integrity. To be able to assess energy needs and determine correct allocation of grid resources, governments need a reliable mapping of distributed PV cells. Because PV's are often privately owned and historical data is unreliable, traditional data collection methods have failed to provide an accurate map of the PV landscape. Furthermore, not all installed PV panels are accurately registered and not all records are up to date. This can result in issues for the renewable energy market as operators need to be able to predict total rooftop solar PV generation over numerous areas, among other concerns. However, the usage of satellite imagery with deep learning can be used as fruitful tool to be able to identify solar PV's to overcome this problem.

Problem Statement

We aim to solve two problems: (a) PV classification - a binary classification task predicting if an image contains any solar panels and (b) PV segmentation - generating pixel masks for the areas in an image that contain solar panels. The inputs to both models were 224x224 RGB images. After being passed through a sigmoid activation function in the final layer, the output for the classifier is the probability that an image contains a solar panel. The output for the segmenter is a binary mask of the input image with 1s indicating the presence of a panel (see below).

(a) Satillite Image (b) Corresponding Mask

The metrics we used for classification were AUC-ROC Curve and F1 score. The ROC is a probability curve, while the AUC indicates the measure of separability - in other words how capable the model is at distinguishing between classes. For segmentation we used Intersection over union. IoU measures the overlap between two bounding boxes or masks.

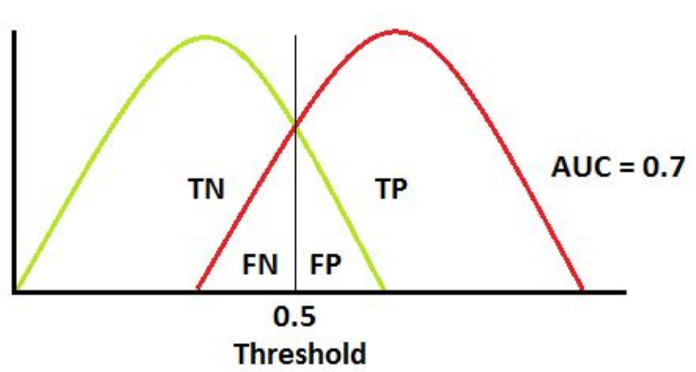

$$F1 = \frac{2 * \text{precision} * \text{recall}}{\text{precision} + \text{recall}}$$

$$\textbf{IoU} = \frac{\text{Area of overlap}}{\text{Area of union}}$$

Dataset

The data set consists of 601 5000-by-5000 pixel TIF ariel images and geospacial coordinates for over 19,000 solar panels. Images span multiple cities in California (Fresno, Stockton, Oxnard and Modesto) and contain a diversity of landscapes including urban, suburban and rural regions. All images have spacial resolution higher than .3m

Methods

For PV classification we used a 34-layer residual network (ResNet-34) with a sigmoid activation function added as the final layer. The network was trained using a Binary Cross-Entropy loss function and Adam Optimizer. The network was pre trained on imagenet for weight initialization.

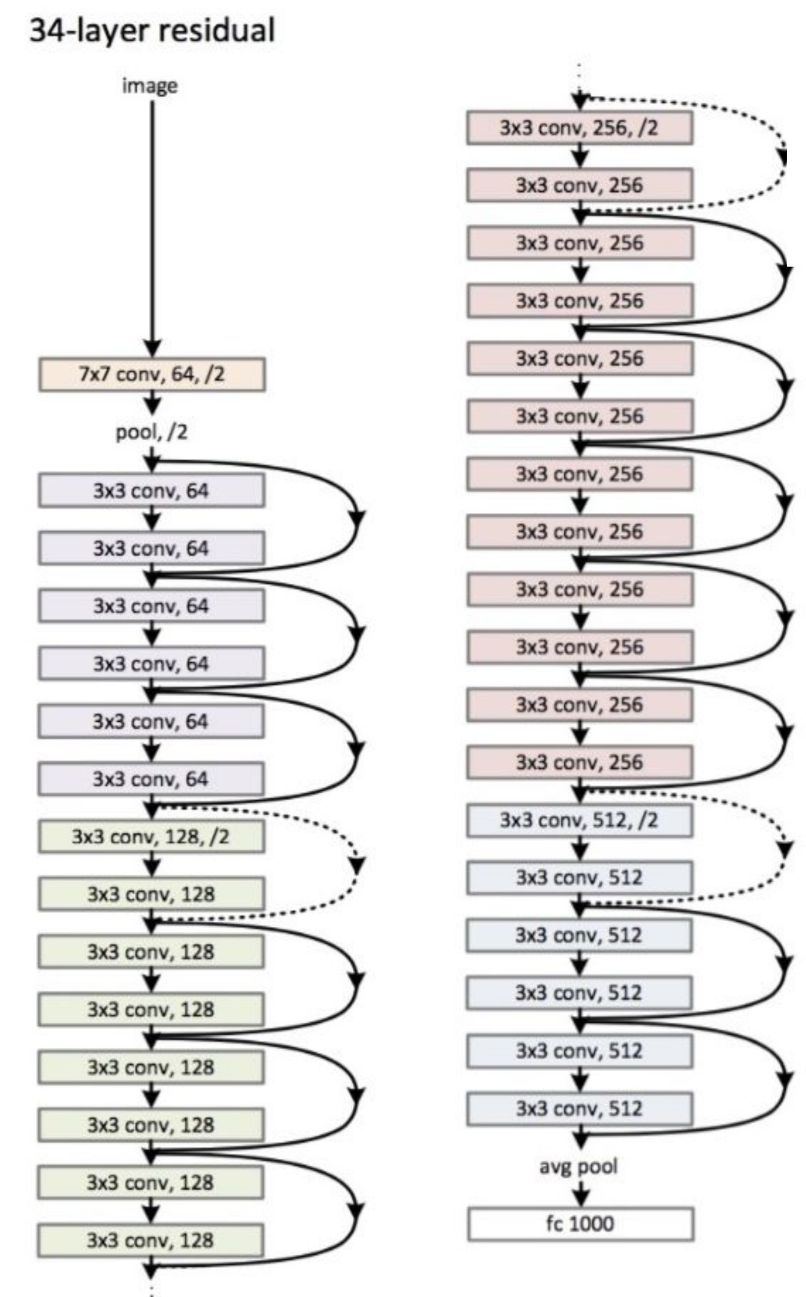

Loss Function: Binary Cross-Entropy

$$\text{Loss} = -(y\log(p) + (1-y)\log(1-p))$$

Optimizer: Adam

$$m_t = \beta_1 m_{t-1} + (1-\beta_1)g_t$$

$$v_t = \beta_2 v_{t-1} + (1-\beta_2)g_t^2$$

$$\hat{m}_t = \frac{m_t}{1-\beta_1}$$

$$\hat{v}_t = \frac{v_t}{1-\beta_2}$$

$$w_t = w_{t-1} - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}}$$

For PV segmentation, we used a UNet architecture. The network consists of a contracting path (left side) and an expansive path (right side). The network was trained using Dice-BCE Loss with an Adam optimizer.

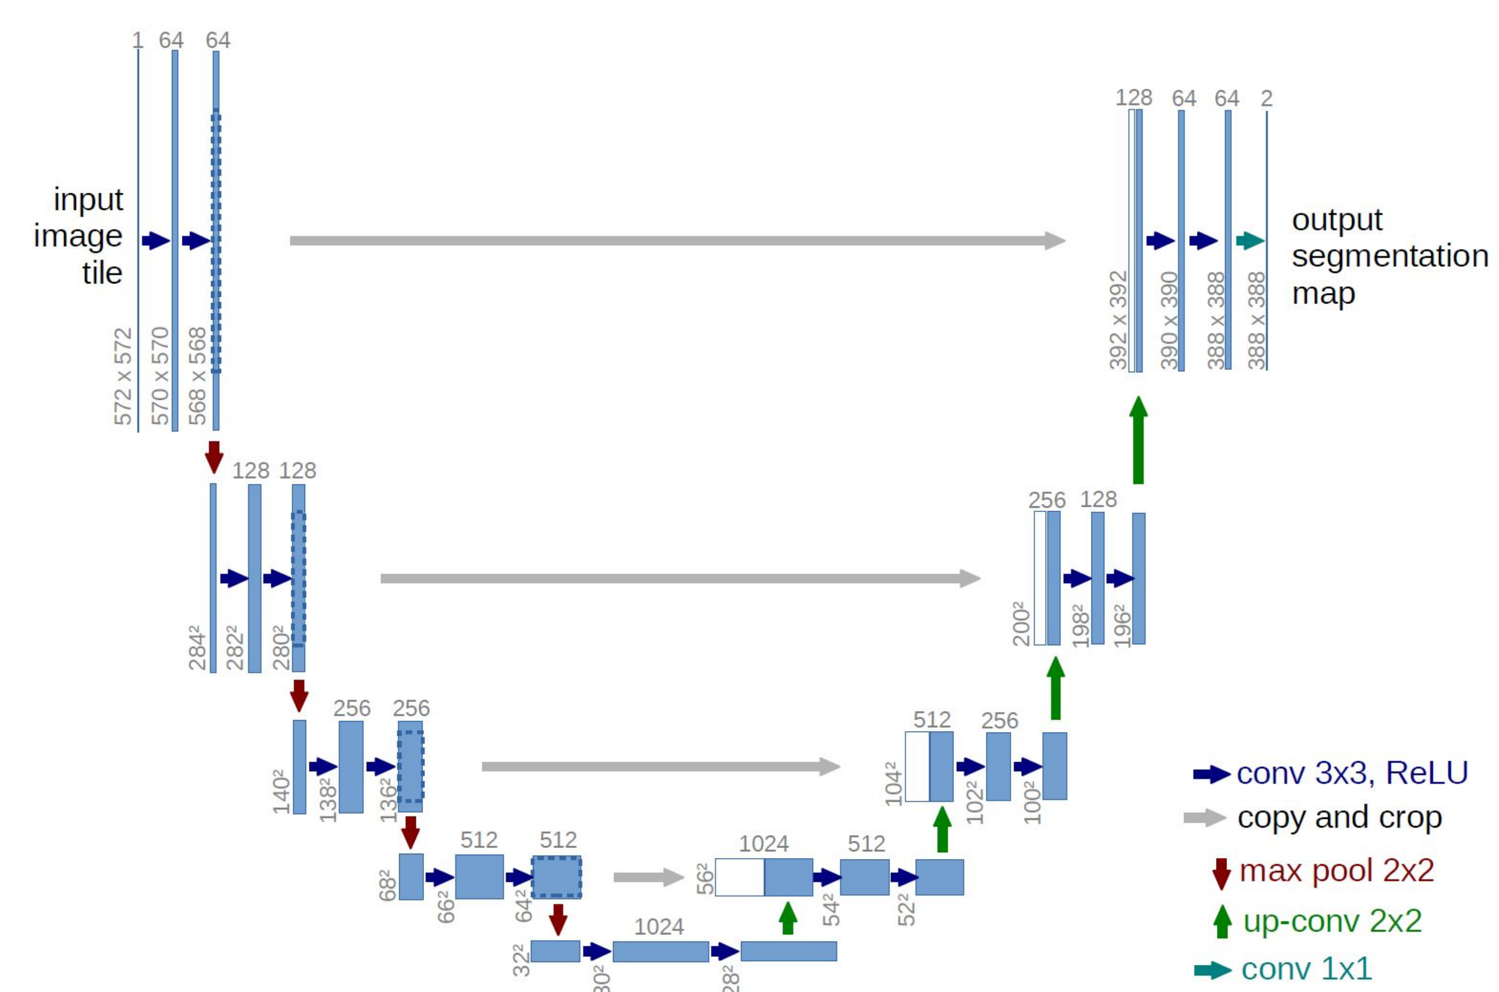

Results

We tuned the hyper-parameters below to achieve a maximum F1 score of .95 and AUC-ROC of .99 for classification and a IoU of .8 for segmentation.

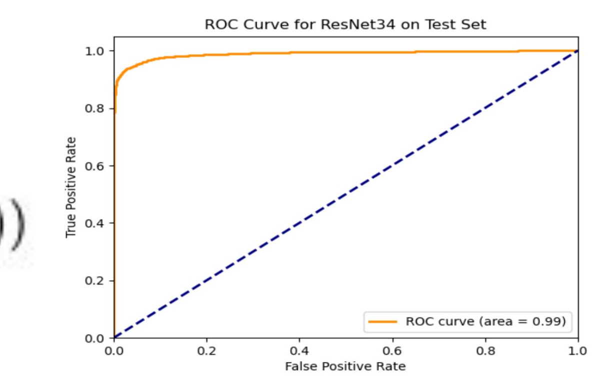

Segmentation Loss Function	Optimizer	IoU
BCE	Adam	0.7552
Dice Loss	Adam	0.7626
BCE-Dice Loss	Adam	0.8049
IoU Loss	Adam	0.7518
Focal Loss	Adam	0.7552

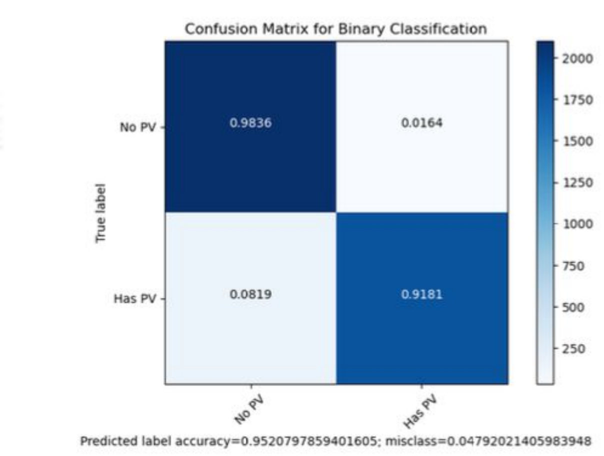

Figure 6. Confusion matrix at threshold = .5 for ResNet-34

Segmen. Optimizer	Segmen LR	IoU
Adam	0.005	0.7317
Adam	0.001	0.789
Adam	0.0005	0.7768

Example of good segmentation, IoU = .98

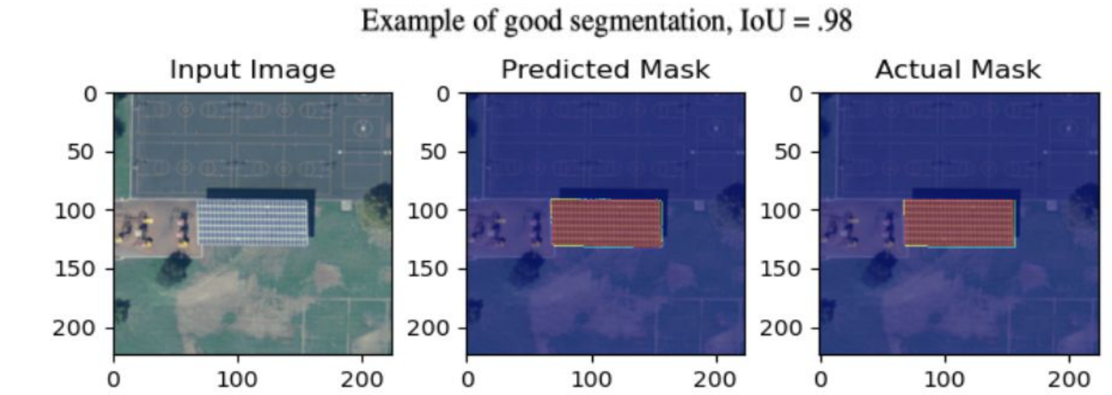

Example of poor segmentation, IoU = .055

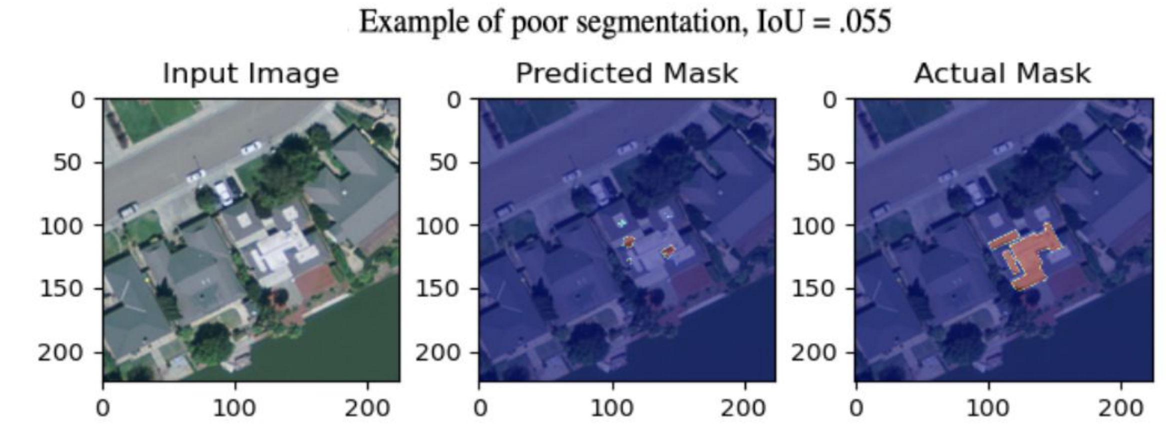

Future Directions

We suggest other work explore PV classification different architectures such as ResNet-50 since it did have a better F1 score than ResNet-34. Additionally, trying segmentation architectures as a Joint Pyramid Upsampling module, two-stream Gated Shape CNN, and an Atrous CNN. Also experimenting with a larger dataset containing images from more than just California to help generalize better to the entire US.