




A CNN-RNN Architecture for Multilabel Classification of the Amazon

Jim Andress and Cristian Zanoci, Stanford University



Motivation

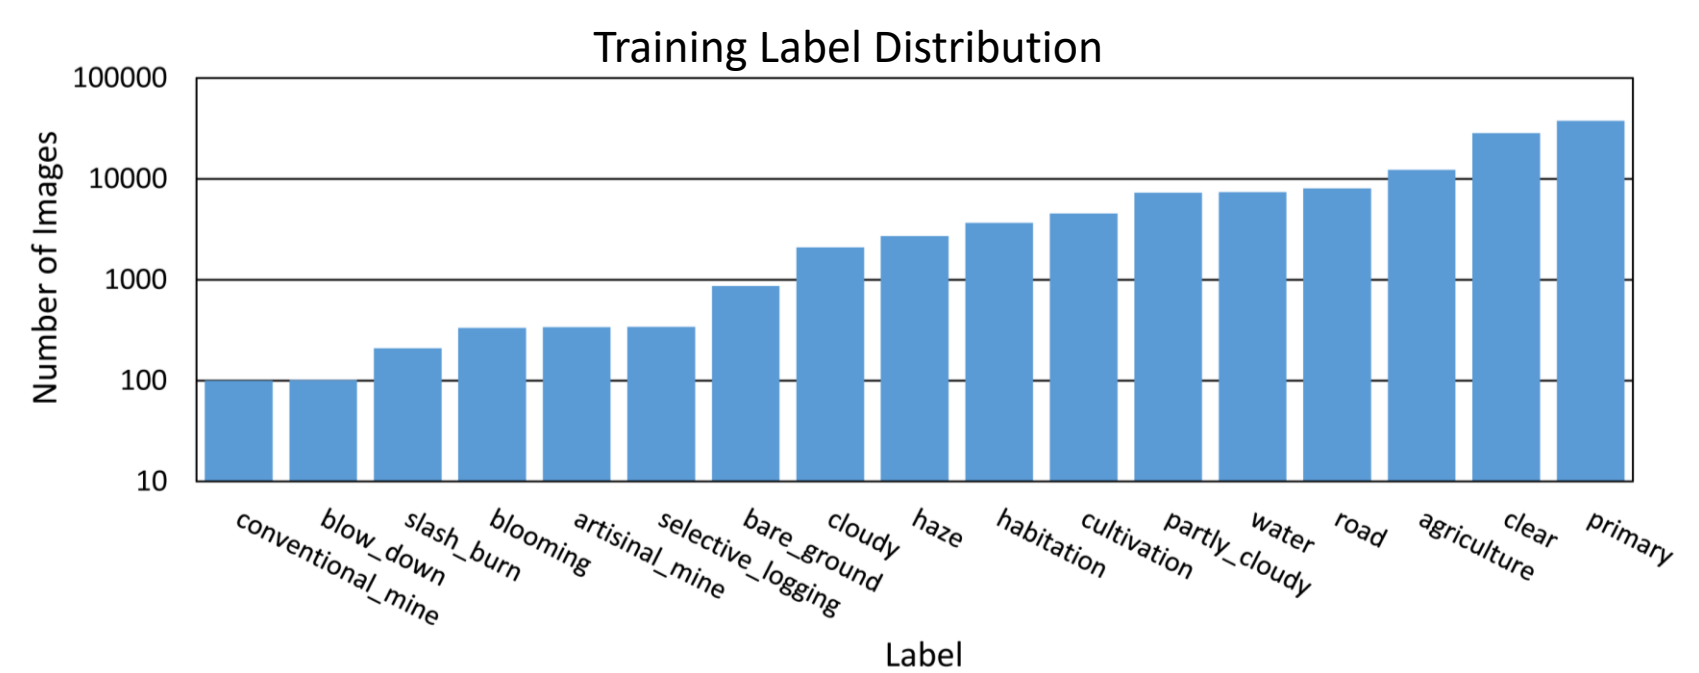
- Deforestation in the Amazon Basin contributes to climate change and reduces biodiversity.
- The size of the Amazon Basin makes human detection of deforestation largely impractical.
- CNN models have proven very successful in classifying satellite images.

Problem Statement

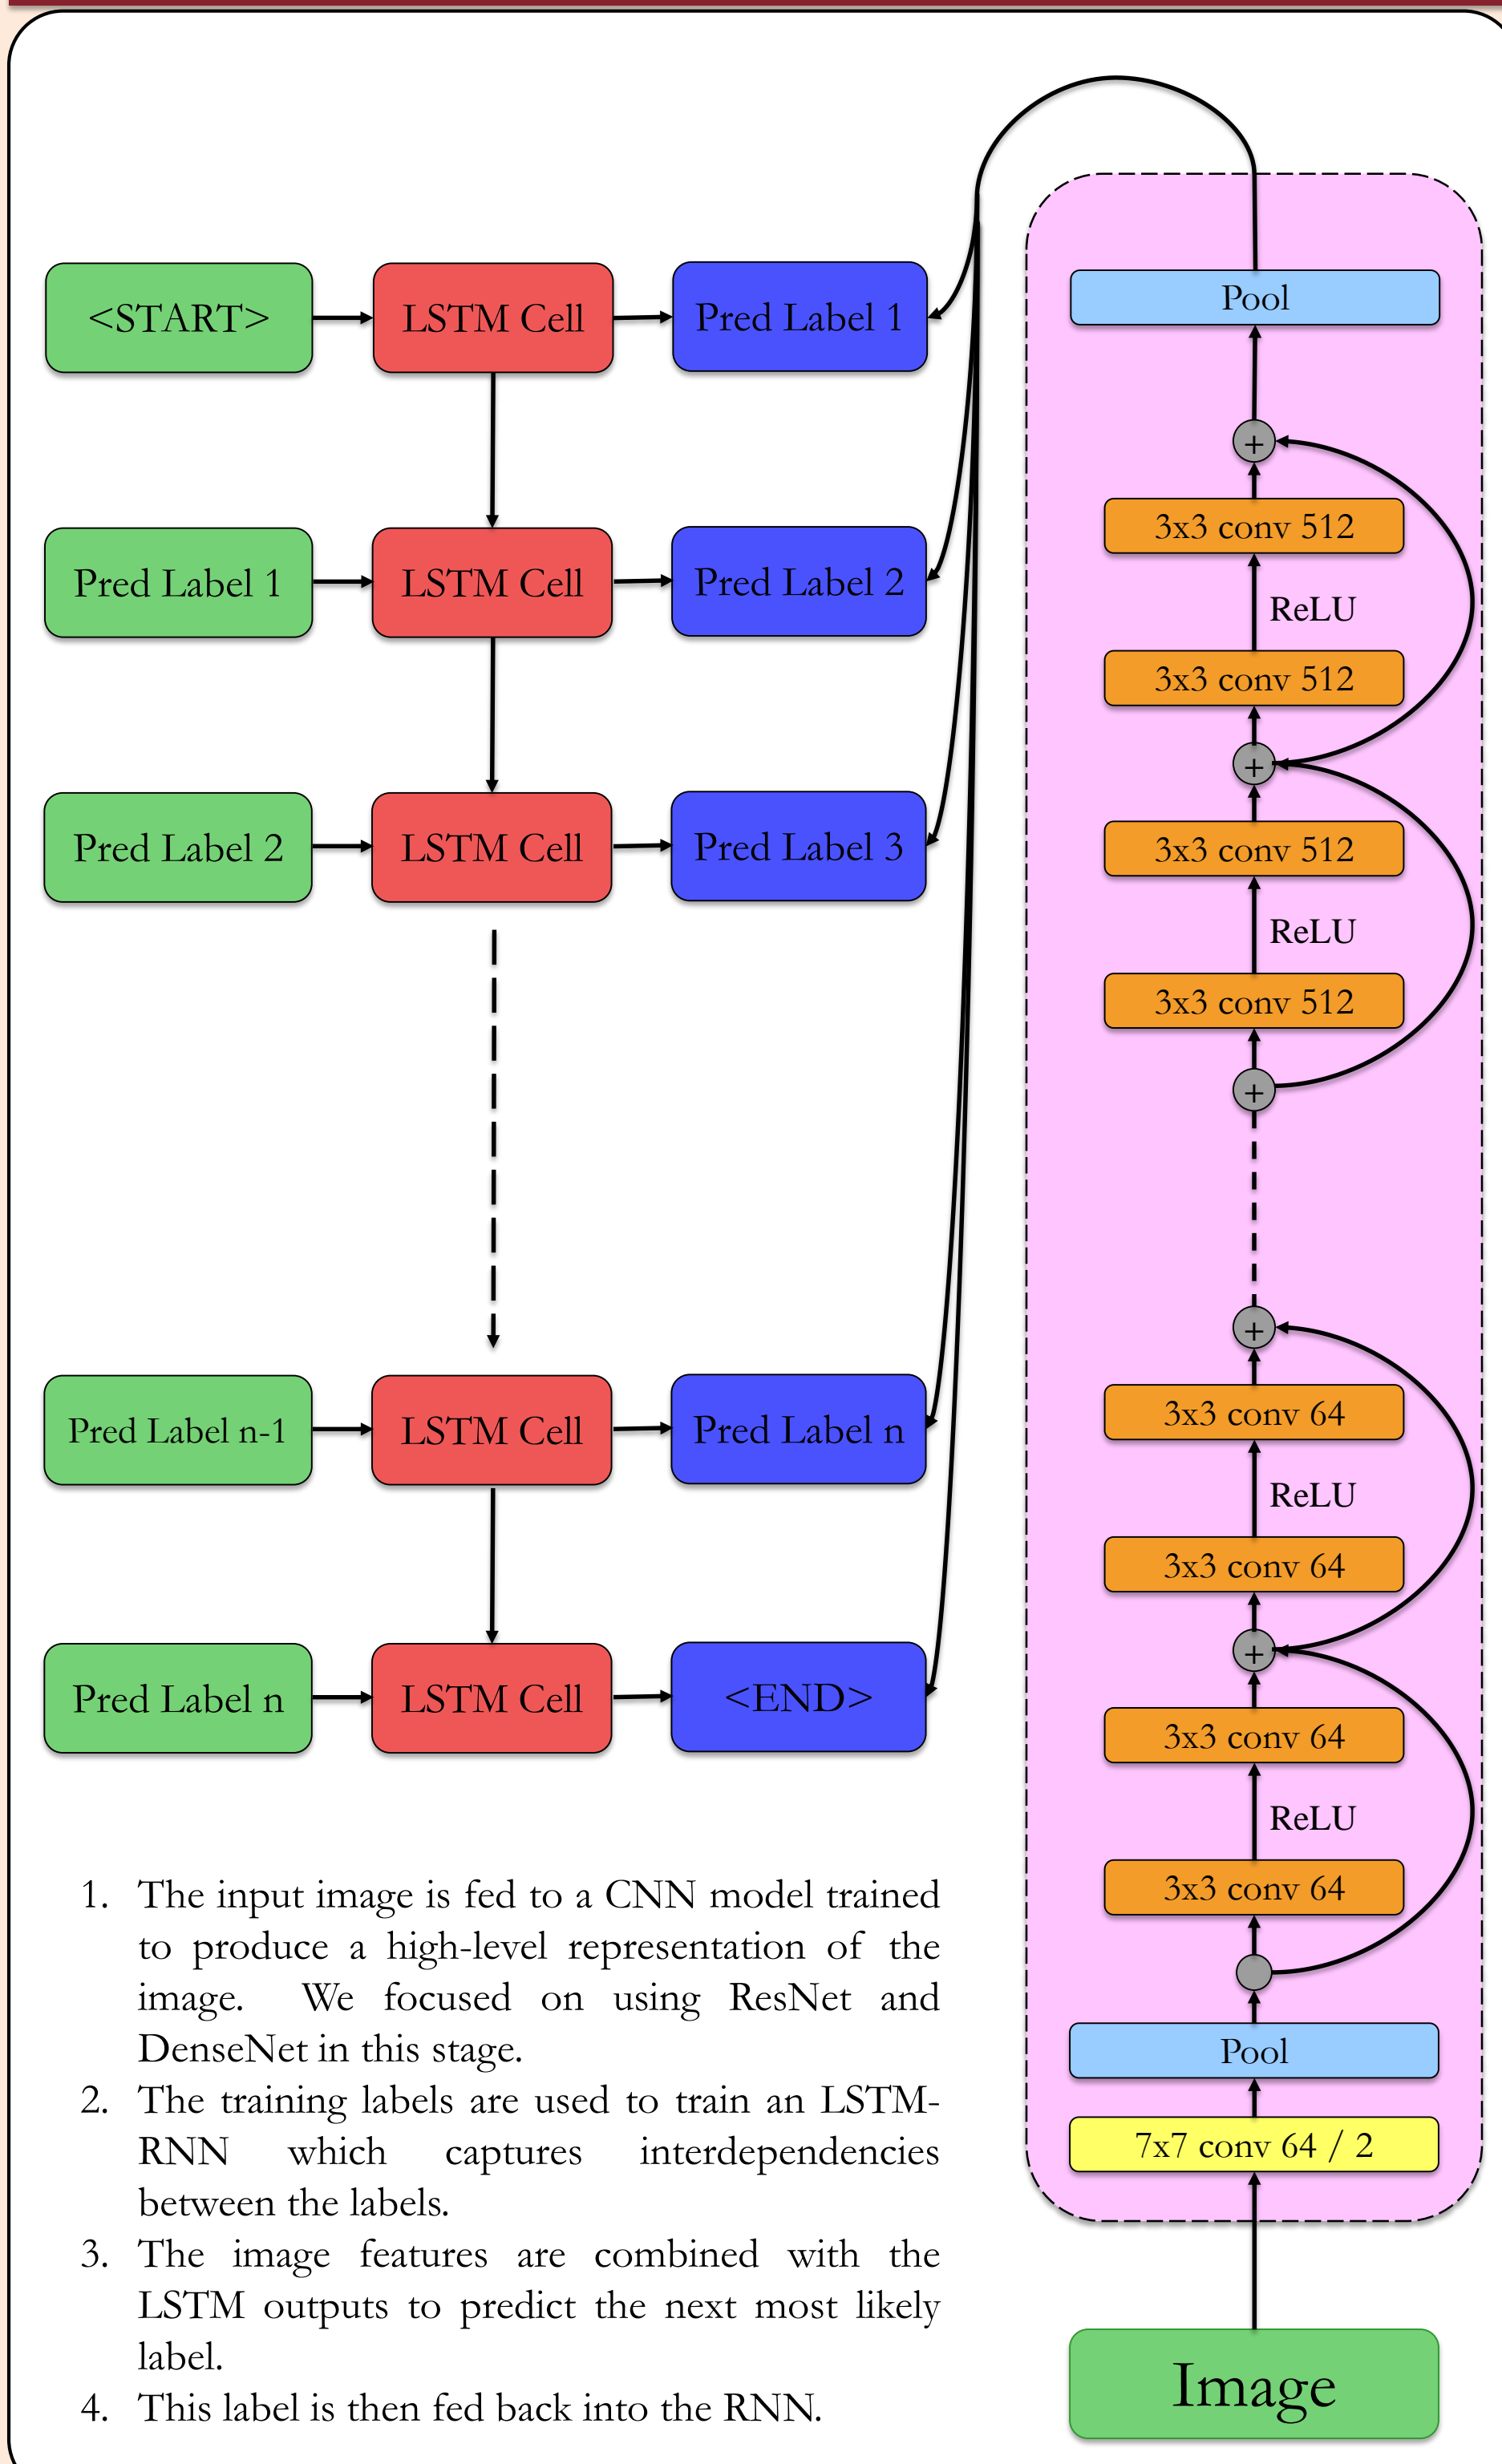
- Input: An image of a small section of the Amazon
- Output: A set of labels describing the image
- Apply a pre-trained CNN model to produce a high-level representation of the image patch.
- Feed this representation to an RNN and iteratively select labels which correspond to the image features.
- Evaluate using the average F2 score of examples.

Planet's Amazon Image Dataset

- 40,480 training images and 61,192 testing images
- Each image is 256x256 pixels in JPEG format
- Also included 4-channel TIFF images, but data/label mismatches made these unusable.



CNN-RNN Architecture



Results

Model	Train F2	Val F2	Test F2
ResNet 18	0.89854	0.91104	0.92229
ResNet 50	0.90589	0.91462	0.92335
DenseNet 169	0.91200	0.92755	0.92276
Ensemble	-	-	0.9255
CNN-RNN	-	0.67569	-
Best on Kaggle	-	-	0.93296

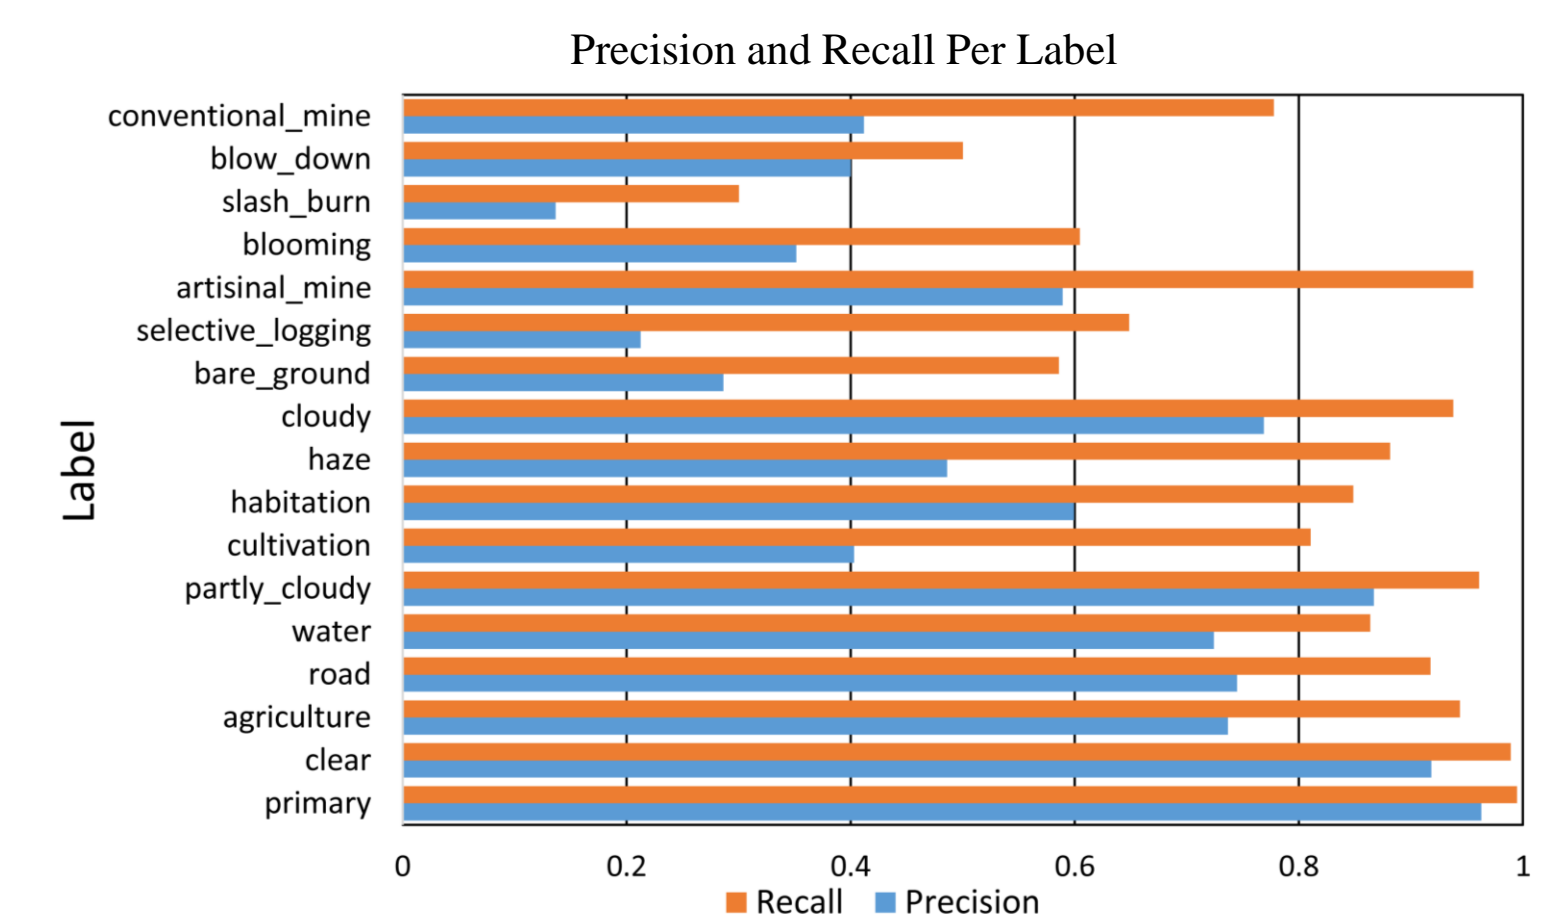


Image Labeling Examples

Predicted: primary, agriculture, clear, cultivation, habitation, road
True: primary, agriculture, clear, cultivation, habitation, road

Predicted: primary, partly_cloudy, water, road
True: primary, partly_cloudy, water

Predicted: primary, road, agriculture, partly_cloudy
True: primary, road, agriculture, partly_cloudy, clear, water

Predicted: primary, clear, road, agriculture, habitation, conventional_mine, bare_ground
True: primary, clear, road, agriculture