Employing a Deep CNN to Track Infection Risk for Schistosomiasis from Satellite Imagery

Zac Espinosa, Ben Gaiarin, Michael Vobejda

Stanford University, CS230 Project, Winter 2018

Fighting a Parasite with Deep Learning

Schistosomiasis is a parasitic disease that, in 2016 alone, infected over 206.4 million people. This debilitating disease is caused by waterborne trematodes known as S*chistosomes* which must live part of their lives in *Bulinus* and *Biomphalaria* freshwater snails before infecting human hosts. In highly endemic areas of disease like the Senegal River Basin (SRB), these snails generally live in patches of vegetation along the shore or floating in the river. Effectively targeting and controlling these snail populations can interrupt the disease cycle. Recent advances in high resolution satellite imagery have allowed for us to see where these sources of snail habitat are, but analysis of these areas by hand is costly and time intensive.

We have built a deep convolutional neural network that classifies and quantifies the prevalence of two types of vegetation–floating ('in the water') and emergent ('near shore')–that provide habitat to the schistosome-carrying snails from high resolution satellite imagery.

Creating a Training Set

- Parsed large-scale, 3m resolution satellite imagery of various SRB regions of interest into smaller 150x150 TIFs. Saved TIFs with infrared band highlighted to show photosynthetic vegetation in red.
- Manually classified 3918 images with two rankings, one ranking per vegetation type (X = floating or emergent), on ranking scale:

1	5	9
No vegetation of type X.	40-50% of shoreline covered with type X.	90-100% shoreline coverage.

Emergent

Emergent: 9
Floating: 1

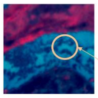

Floating

Emergent: 9
Floating: 8

Emergent: 1
Floating: 1

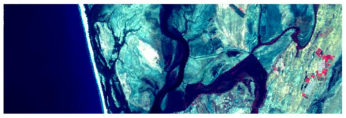

Example of an original file image.

All images courtesy of Planet Labs Inc.

Iterating through Designs

#	Architecture/Parameter Change	Results
1	CONV2D → RELU → MAXPOOL → CONV2D → RELU → Z3 (BASE = CONV2D → RELU → MAXPOOL → CONV2D → RELU) Z3 → $Z3_{emergent}$ → BASE → MAXPOOL → FC Z3 → $Z3_{floating}$ → BASE → MAXPOOL → FC ○ Learning rate = 0.009 ○ Train = 95%, Dev = 5% ○ Mini batch size = 2 ○ CONV stride = [1,1,1,1] ○ MAXPOOL stride = [1,8,8,1] ○ NOT run with data augmentation. ○ NOT run with *weighted* softmax cross entropy.	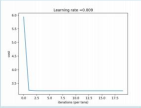 Train Floating: 0.847 Test Floating: 0.875 Train Emergent: 0.743 Test Emergent: 0.833 *Accuracies are artificially high
2	○ Learning rate = 0.001 ○ Weight = 100.0 applied to class 1. ○ Fixed matrix bug; cost function corrected. ○ Accuracies still artificially high, remain same. ○ Applied six methods of data augmentation to all images not of rank 1-1. ○ NOT run with *weighted* softmax cross entropy.	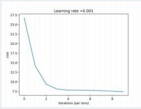
3	CONV2D → RELU → MAXPOOL → B(TANH)*3 → Z5 (B(RELU) = CONV2D → RELU → MAXPOOL) Z5 → $Z5_{emergent}$ → TANH → FC Z5 → $Z5_{floating}$ → TANH → FC ○ Mini batch size = 16 ○ Class weights for 1-9: [1.5,1,1,1,1,1,1.1,1.3,1.5] ○ Accuracy calculation fixed to count a predicted rank correct if: Prediction – Correct <= 4 if rank ≠ 1 Prediction – Correct <= 1 otherwise	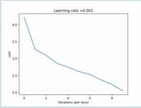 Train Floating: 0.806, Ones: 0.937 Test Floating: 0.722, Ones: 0.963 Train Emergent: 0.656, Ones: 0.921 Test Emergent: 0.724, Ones: 0.968
4	B(RELU) → B(TANH)*3 → Z5 Z5 → $Z5_{emergent}$ CONV2D RELU CONV2D RELU TANH → FC Z5 → $Z5_{floating}$ CONV2D RELU CONV2D RELU TANH → FC	Train Floating: 0.805, Ones: 0.932 Test Floating: 0.722, Ones: 0.953 Train Emergent: 0.653, Ones: 0.928 Test Emergent: 0.719, Ones: 0.978
5	Accuracy-affecting bugs fixed. Bugs included: ○ Images of rank 1-1 augmenting, and not images of rank 2-2. Augmented files not persisting. ○ Lengths of Y_F (floating labels), Y_E (emergent labels), and X (images) not equal. Implemented softmax-less, binarized version of model, classifying 1-1 as 0 and other rankings as 1. Changed rank threshold: ○ Prediction – Correct <= 3 if rank ≠ 1	(Results shown in Final Results section.)

Final Results

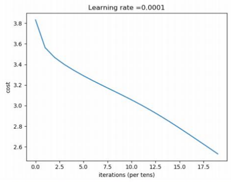

Accuracy Type	Non 1-1 Images	1-1
Train, Floating:	0.763	0.912
Test, Floating:	0.752	0.905
Train, Emergent:	0.677	0.845
Test, Emergent:	0.694	0.889

Reflections on heat maps (below):
- Emergent vegetation classified with good precision, but model generates many false positives.
- More accurately-classified, larger training set necessary for model to learn floating vegetation.

Reflections on cost and accuracy (left) from **multiclass** implementation of **design iteration #5**:
- Model accurately distinguishes between 1-1 and non 1-1 images.
- Model does not generalize well to unlabeled testing data.

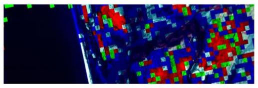

Heat Map A: From multiclass implementation of design iteration #5. Floating = Red, Emergent = Blue, Both = Green.

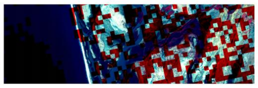

Heat Map B: From binarized implementation of design iteration #5. Emergent = Red.

Implications

Our work suggests that a deep convolutional network can be used to map localized habitat suitability for schistosome-carrying snails. Our model indicates that changes in the prevalence of vegetation types can be monitored from satellite imagery on a scale unmatched by traditional methods. Further refinements to our model are needed to make the network more precise, more accessible for researchers, and able to classify vegetation on smaller scale images.

Future work

- Improving the accuracy and, in particular, the preciseness of the model's predictions by:
  - Using higher resolution (<1m) satellite imagery
  - Comparing satellite imagery with drone imagery during manual labeling of training data.
  - Continuing experimentation with parameters.

Contact

Zac Espinosa, Ben Gaiarin, Michael Vobejda

Produced at Stanford University. Made in collaboration with the De Leo Lab.

With thanks to Columbae for meals and friends.

Email: zespinos@stanford.edu, bgaiarin@stanford.edu, mvobejda@stanford.edu