Plankton Classification with CNNs

Sidney Batchelder

WOODS HOLE
OCEANOGRAPHIC
INSTITUTION

Plankton!

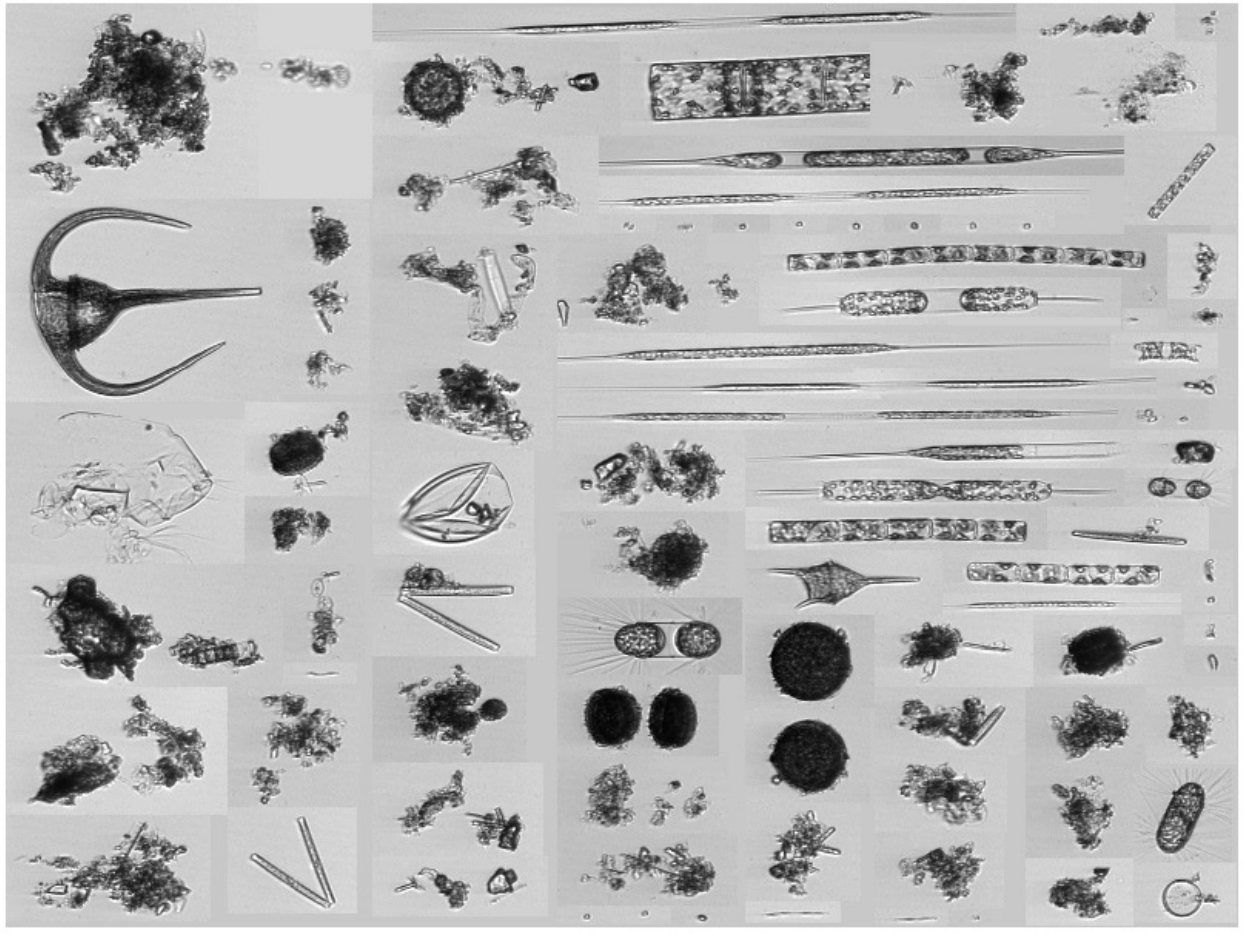

Problem Overview

- IFCB takes 100k+ plankton images per day
- Too many for manual annotation
- Previous feature-based classifier (random forest)
  - ~70% effective
- Let's do better!

Neural Nets

- Emulate, broadly-speaking, the way signals flow through neurons of a biological brain

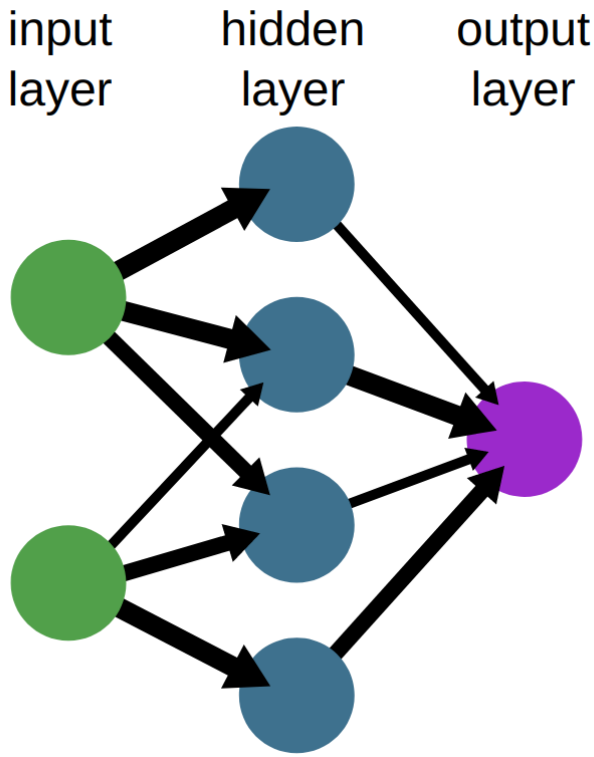

Training NNs (supervised)

- Known inputs-outputs pairs (labeled data)
- Output error (loss)
- Updating node weights (backprop)
  - backward propagation of errors
  - optimizer
- Iterate!
- Don’t Overfit.

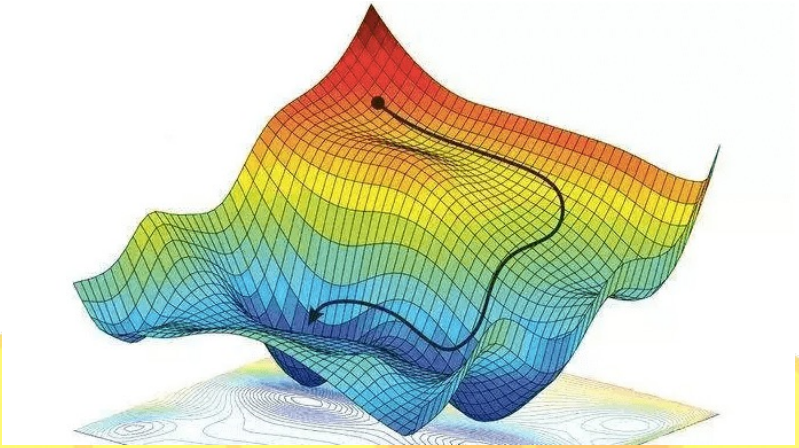

Backprop: it gets complex

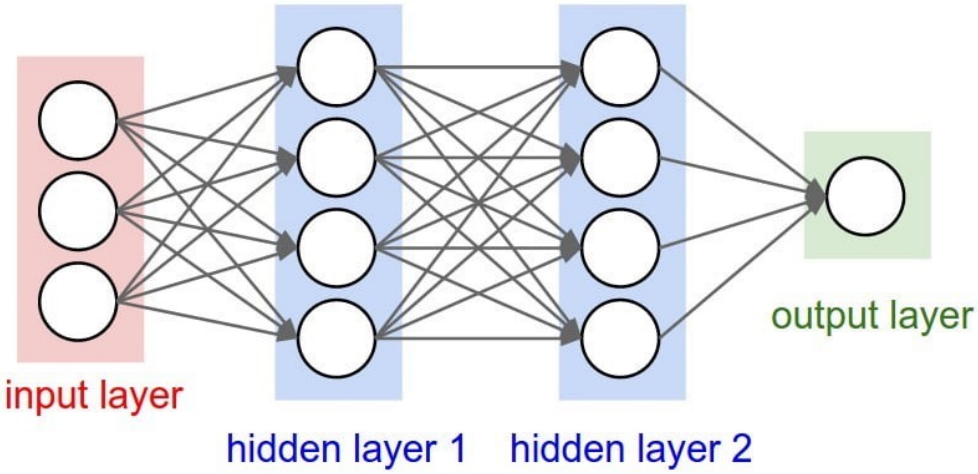

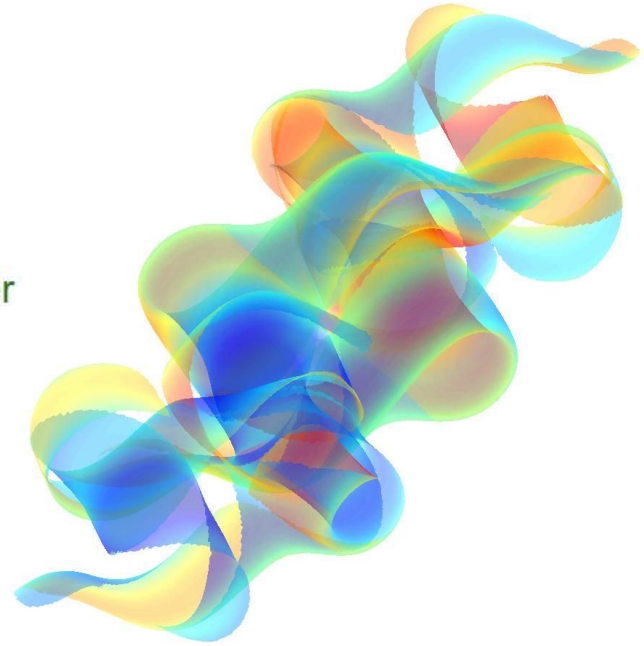

Lots of configurations exist...

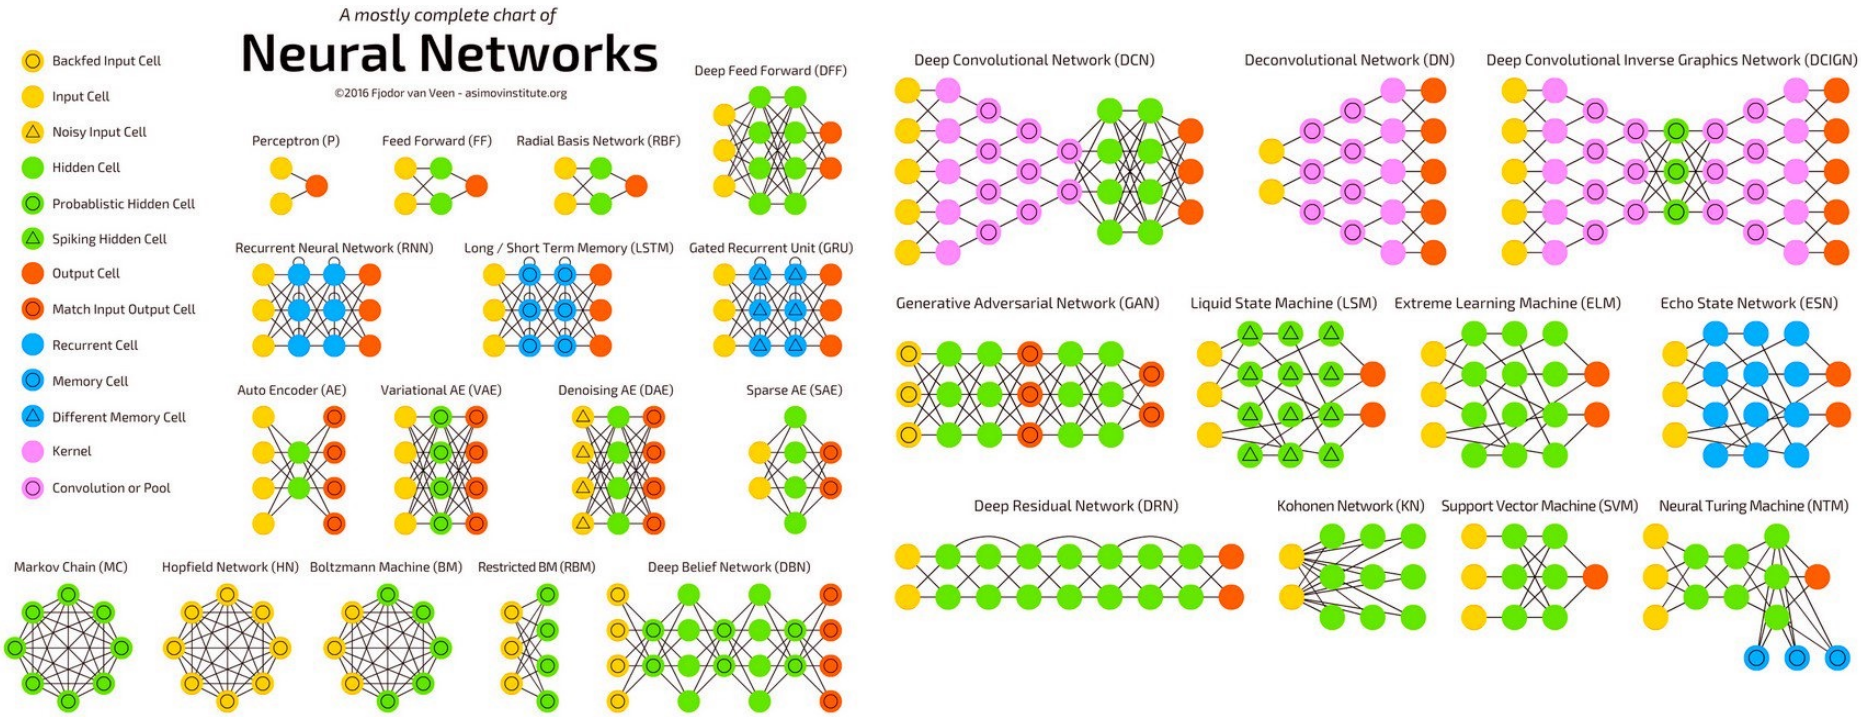

So what are we working with here?IFCB Data

- 300k labeled images over 100-150 classes
- Billion+ images in total
  - MVCO dataset, the largest spanning 17 years
- Need:
  - Efficient processing of images
  - Differentiate output for different classes
  - Enough complexity to handle lots of cases/details

What do we need?

- Deep
  - Many hidden layers
  - Allows for more complex “understanding”
- Convolutional
  - Method for reducing complexity of input for next layer by identifying features, ie patterns.
  - Great for rich media (images, audio, 3D)
- Multi-Label Classification
  - Many nodes in output layer
  - Cross-Entropy loss function

ConvNets

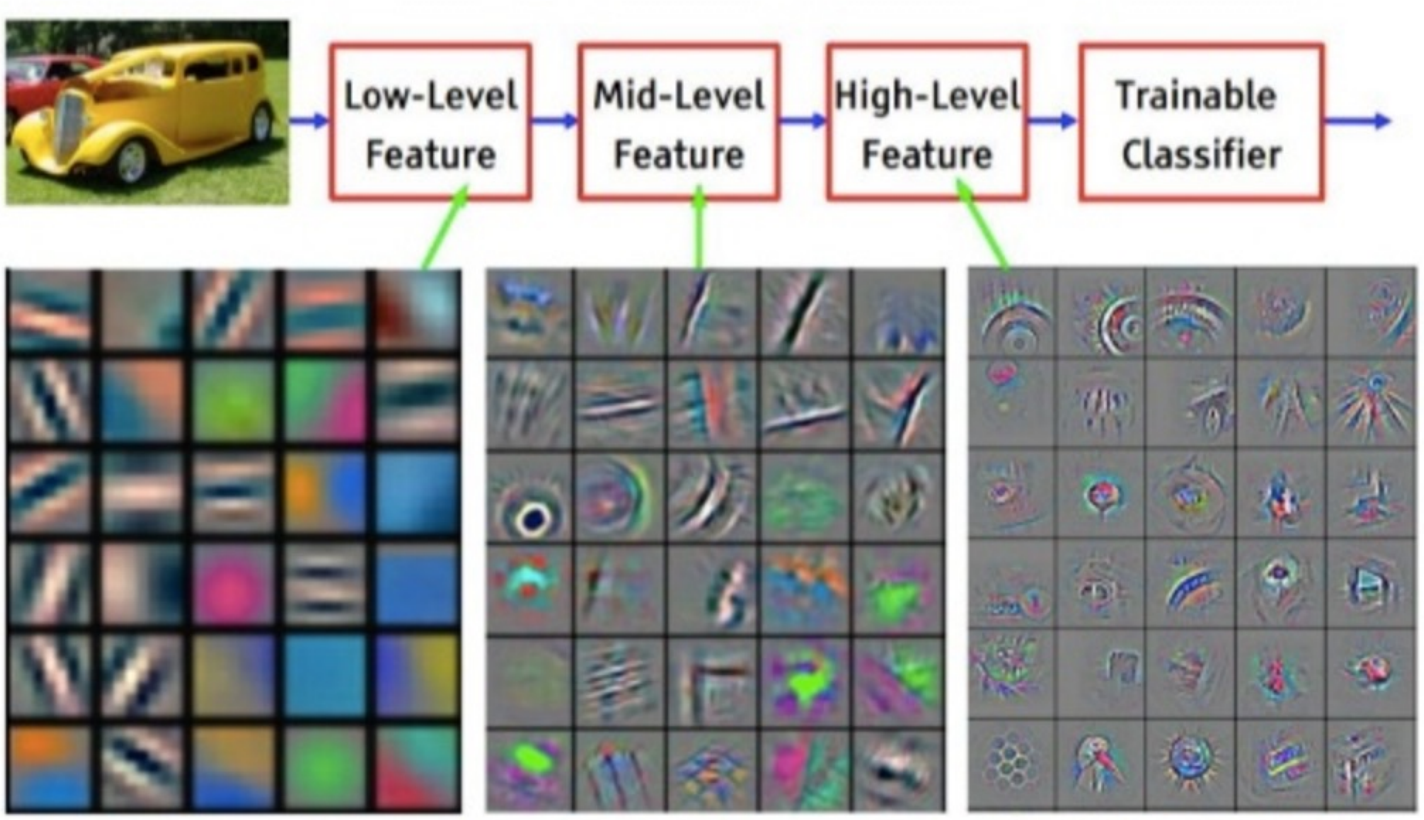
Low-Level Feature
Mid-Level Feature
High-Level Feature
Trainable Classifier

ConvNets (cont…)

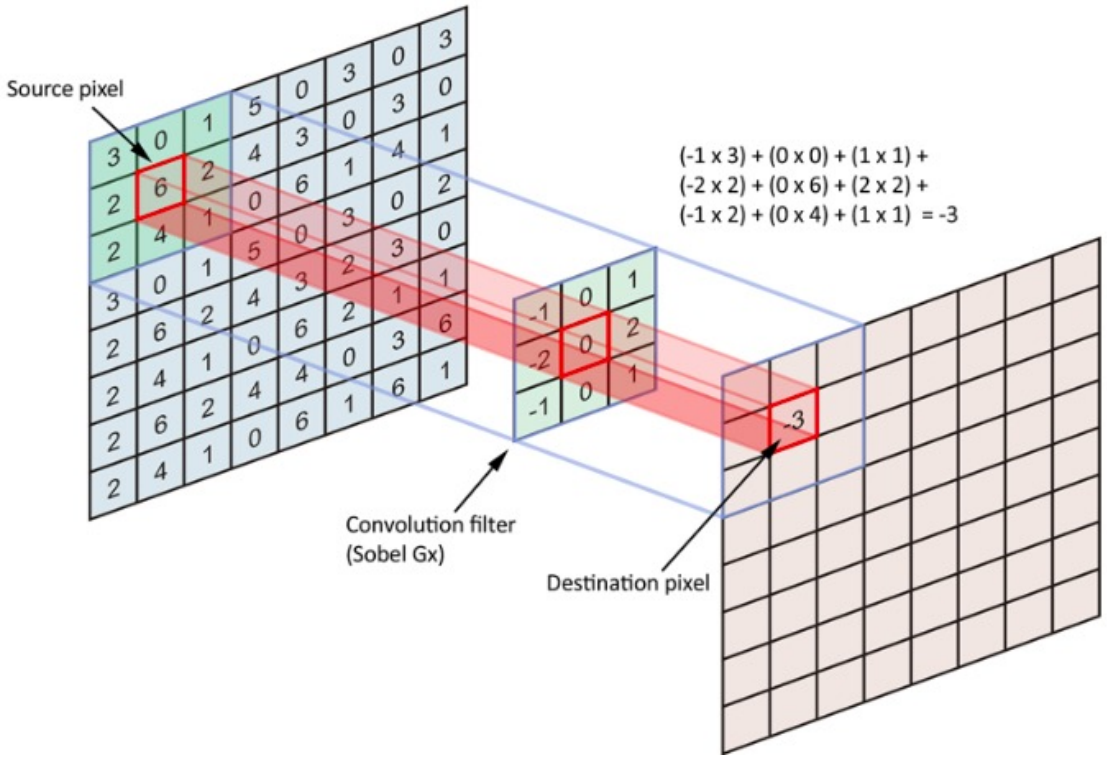

Leveraging Existing Architectures

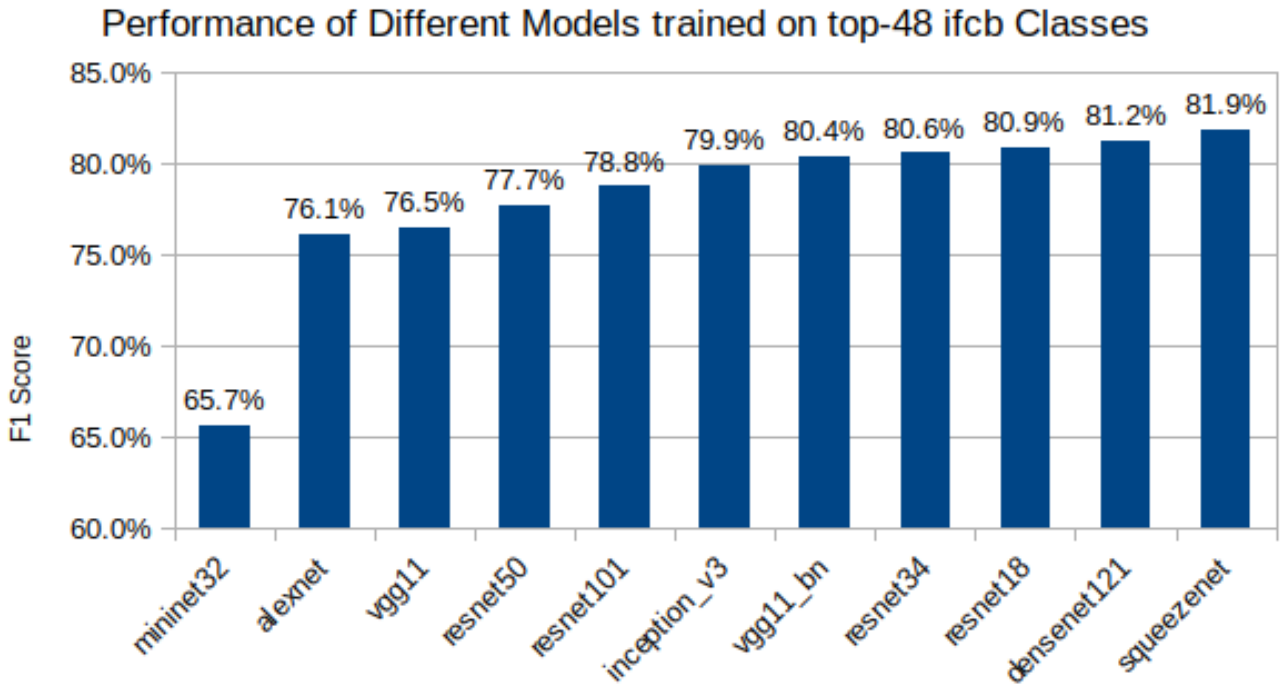

* initial architecture explorations

Inception V3

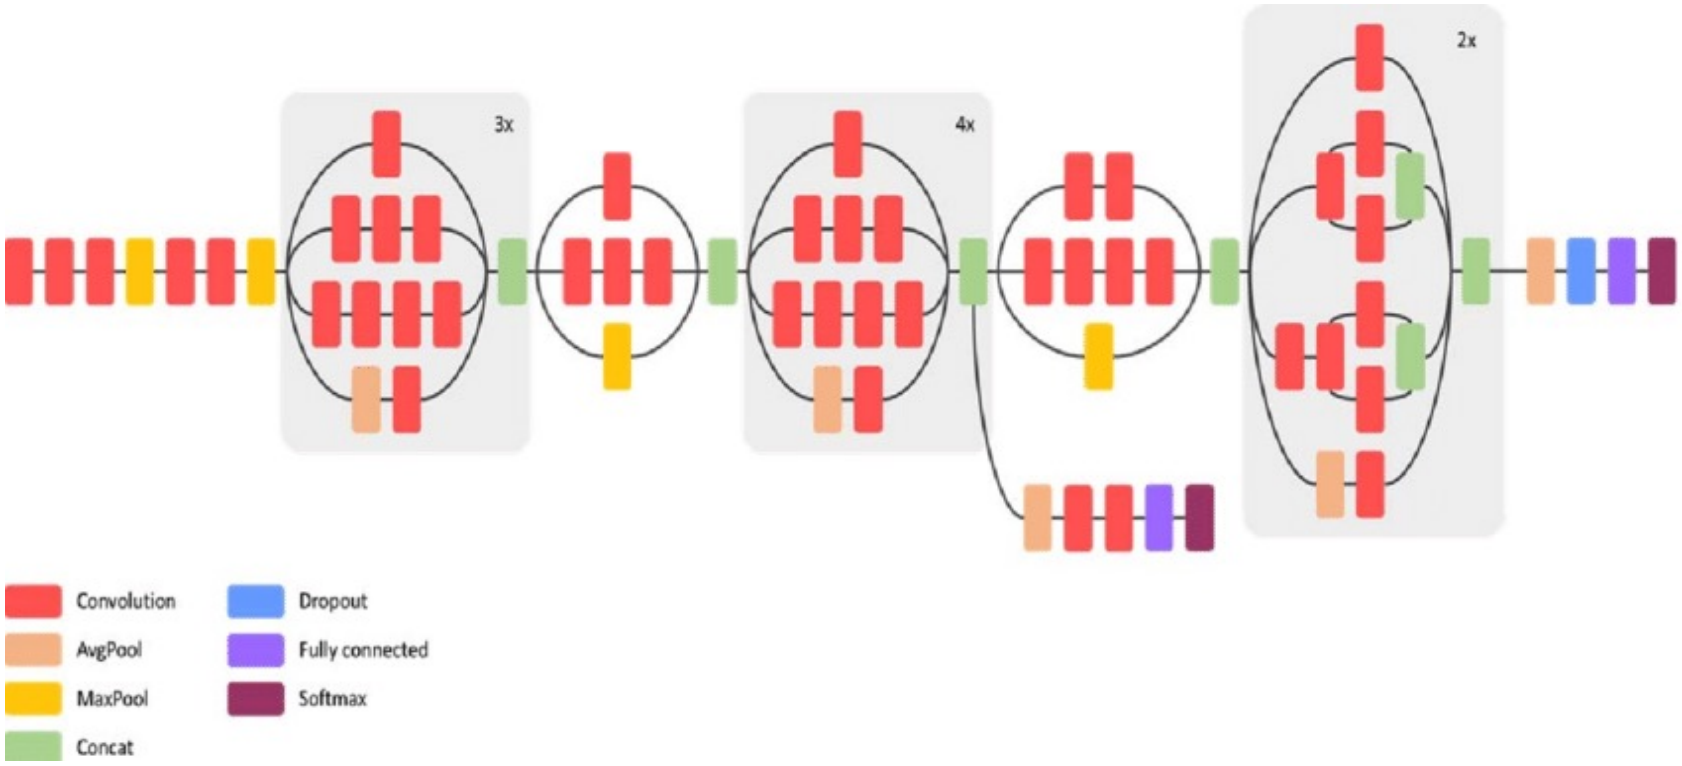

Results: Training

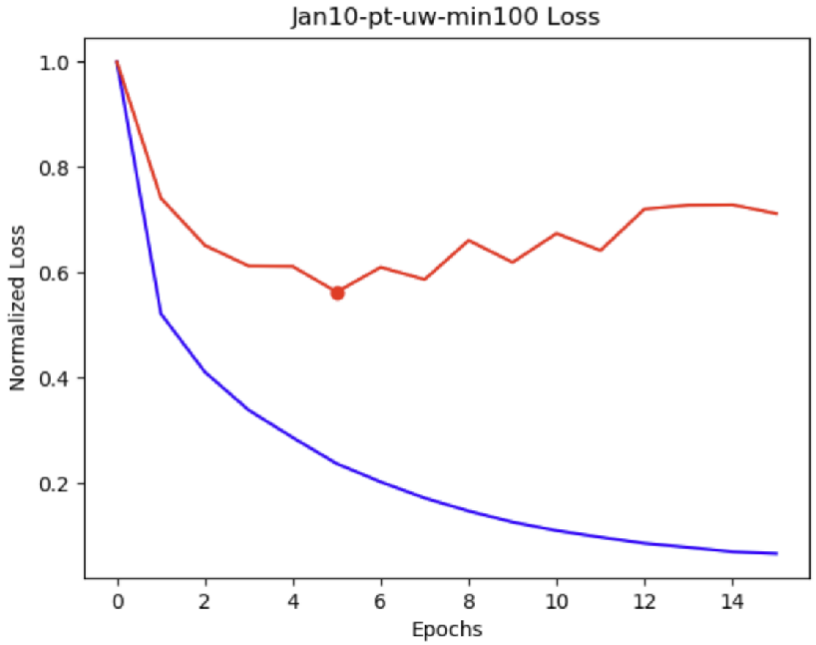

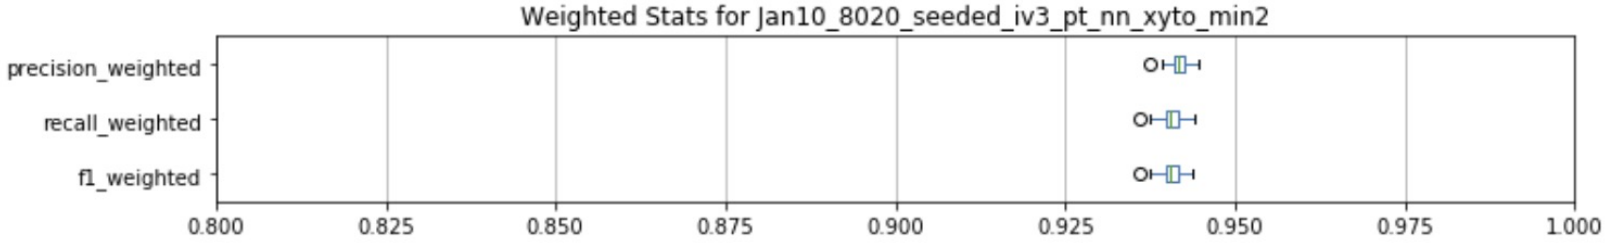

Results: Confusion Matrix

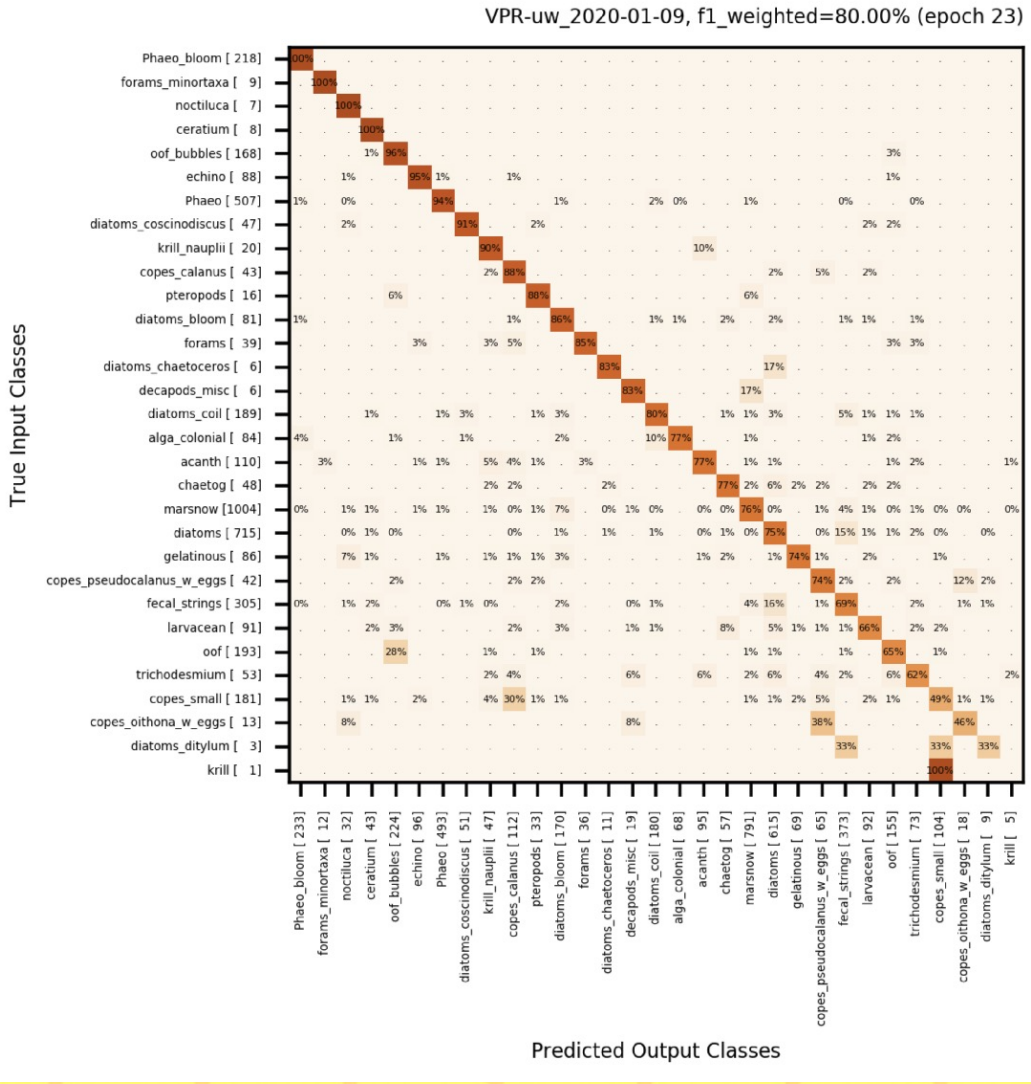

Results: Confidence

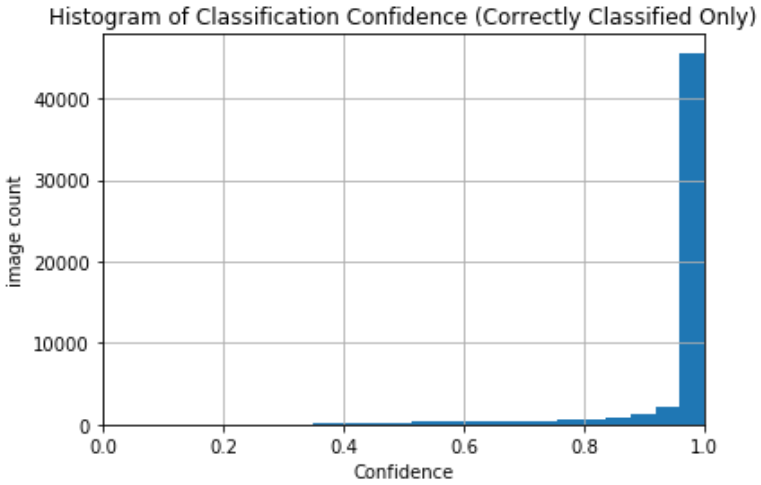

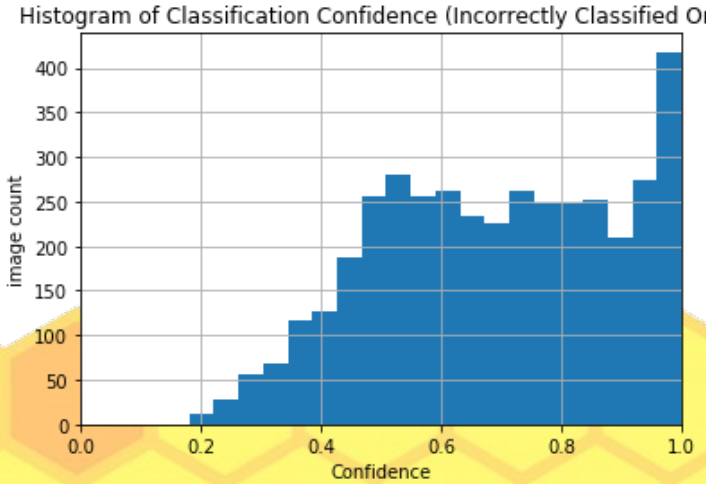

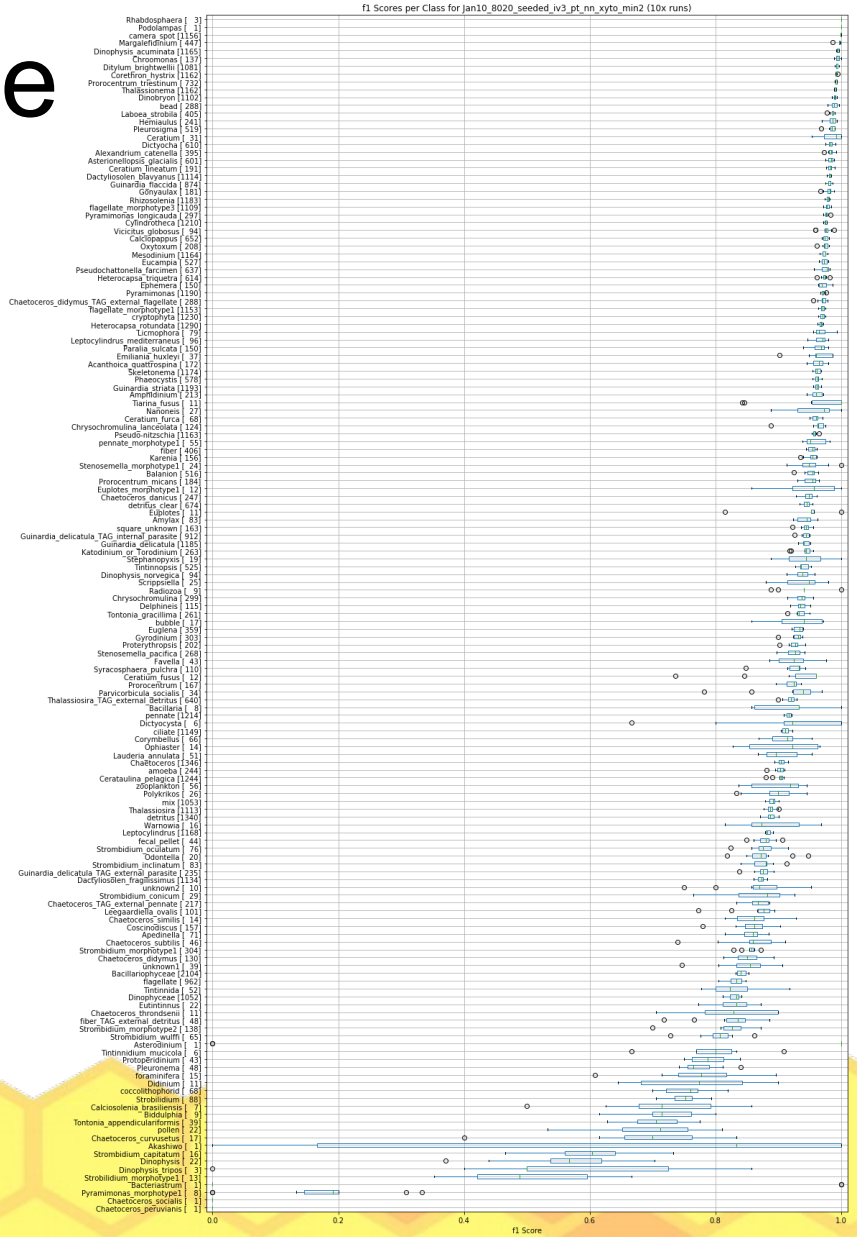

Improving IFCB products

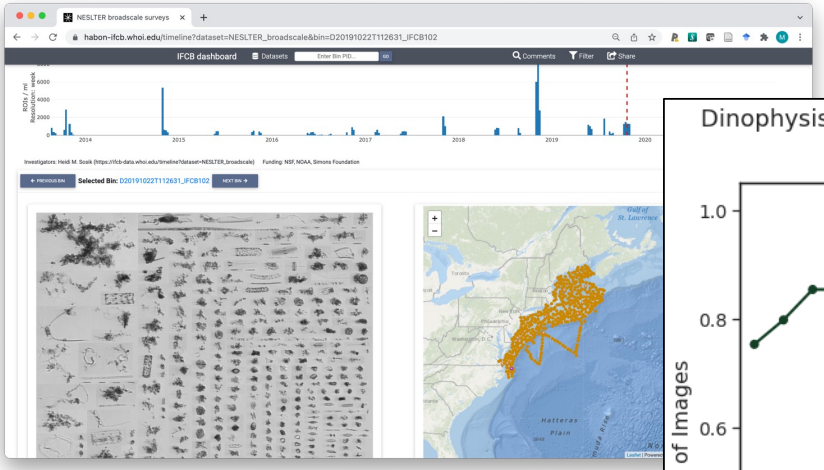

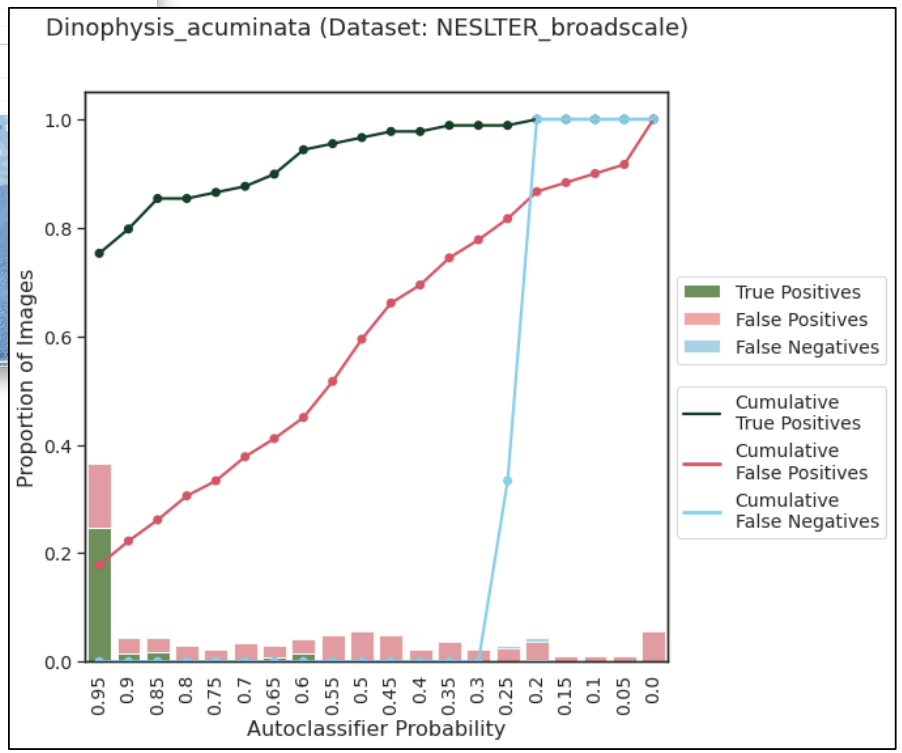

How are winning scores distributed among correct (TP) and wrong (FP) classifications?

Improving IFCB products

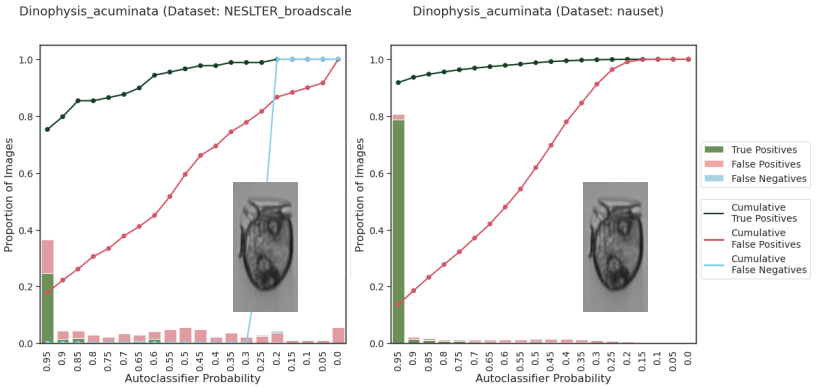

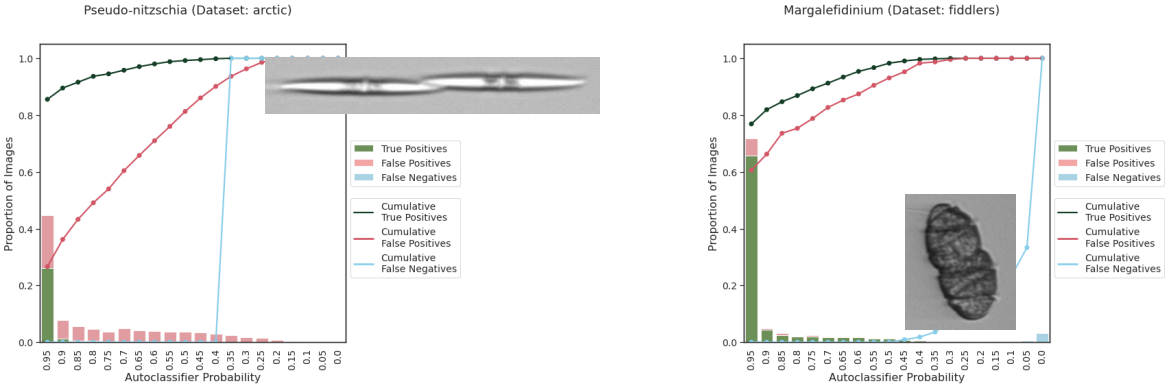

>95% prob threshold produces ‘cleaner’ classifier product than simple winning prob

Details vary by species and time/location

Human annotation of images remains the gold standard

Need a mechanism to push these directly into HAB hub

Results: Class Counts

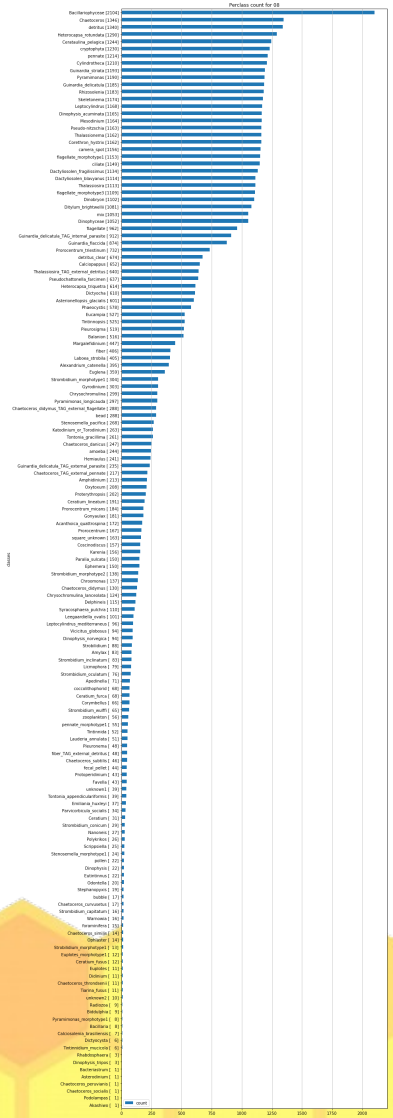

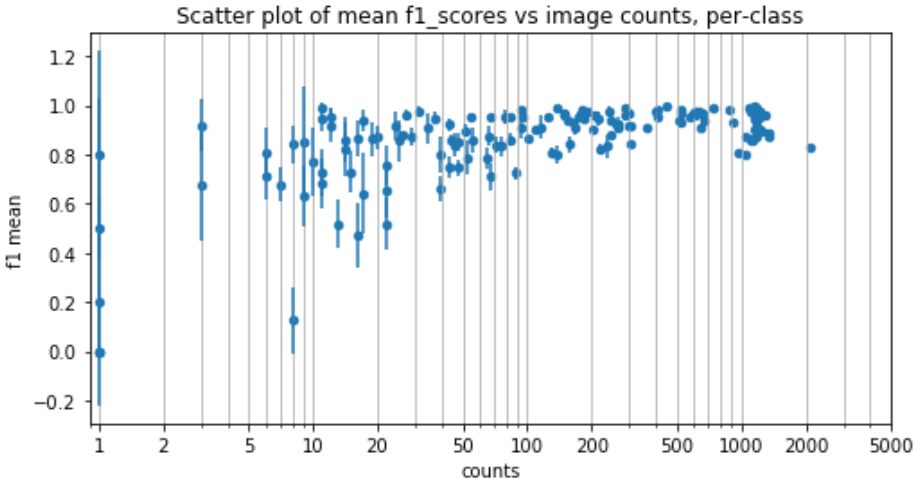

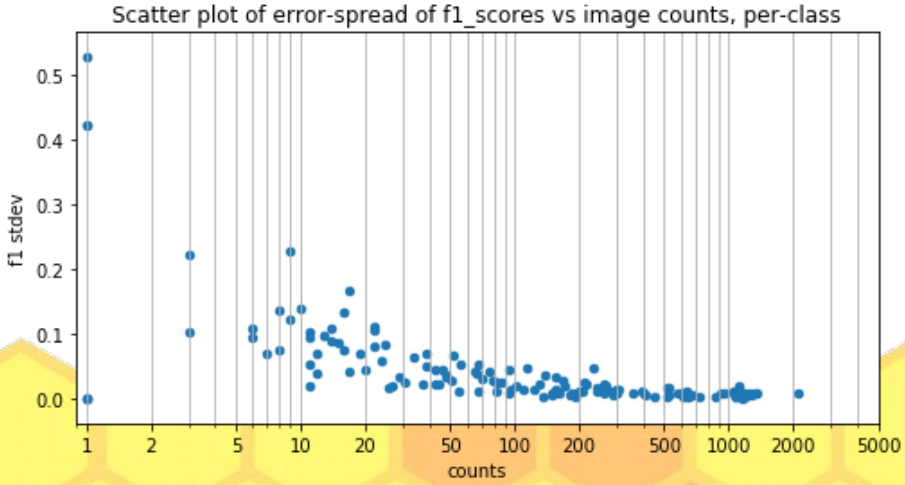

HPC Stats

- GPU Node: 4x NVIDIA Volta 100's
- Training
  - half hour per epoch, ~16 epochs per run typically
  - 8 hours per fully trained model
- **Running**
  - MVCO data-subset: 29 million images
  - Total time: 19.5 hours
  - Single GPU Rate: 400 images/sec
  - Max Rate: 1600 images/sec with 4 GPU's

Outcome

Process 17 years of data
(7.7 TB)

in approx 1 week

at 90%+ effectiveness
with 100+ classes

Classification on the Edge

NEW

- Sending image data by satellite is costly
- Sending counts and statistics is not!
- Nvidia Jetson modules + other hardware
- Realtime processing

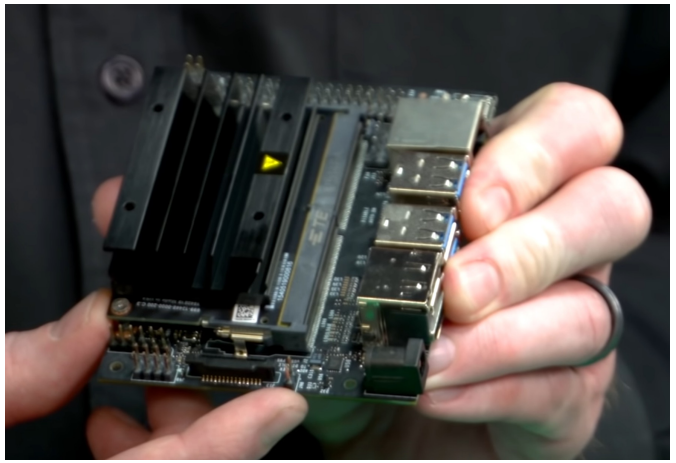

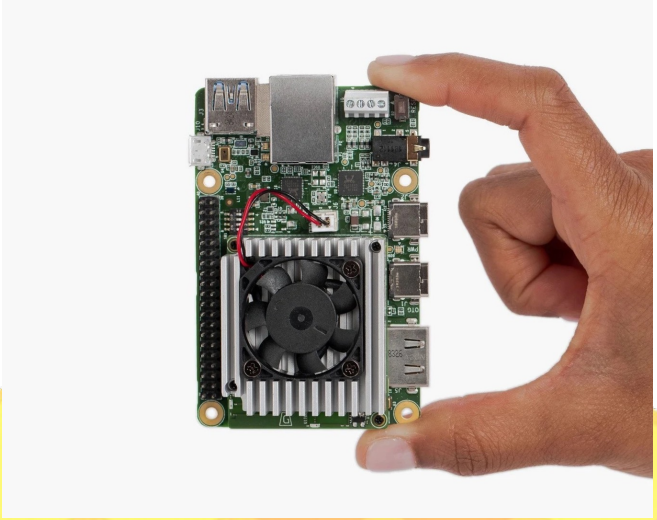

Resources and Obstacles

- GitHub Repo
  - github.com/WHOIGit/ifcb_classifier
- Trained models not currently published/available
- No active support for 3rd parties…

- But you're welcome to fork the project!

Acknowledgements

- Dr. Heidi Sosik
- Members of the Sosik Lab
- Michael Brosnahan
- WHOI IS App Dev Team

Thank you!