Predicting Ocean Health One Plankton at a time

Abhilash Kumar
12014

Peeyush Agarwal
12475

Motivation

Critically important to our ecosystem

- Represent the bottom few levels of food chain
- Play an important role in ocean’s carbon cycle

Population levels are an ideal measure of the health of world’s oceans and ecosystems

Traditional methods are

- Time consuming
- Cannot scale for large-scale studies

Could take a year or more to manually analyze the imagery volume captured in a single day

A better approach :

- Use underwater imagery sensors for capturing images
- Automated image classification using machine learning

Objective

To create an algorithm that given an image, assigns class probabilities for various plankton classes.

Dataset

Provided for Data Science Bowl competition

Contains 121 Classes

Consists of :

- 30,000 labeled images
- 130,000 test images

Challenges

- Many different species with varying size
- Image can have any orientation within 3-D space
- Ocean replete with detritus that have no taxonomic identification
- Sometimes difficult for even experts because of noise
- Presence of "unknown" classes

Methodology

Computer Vision

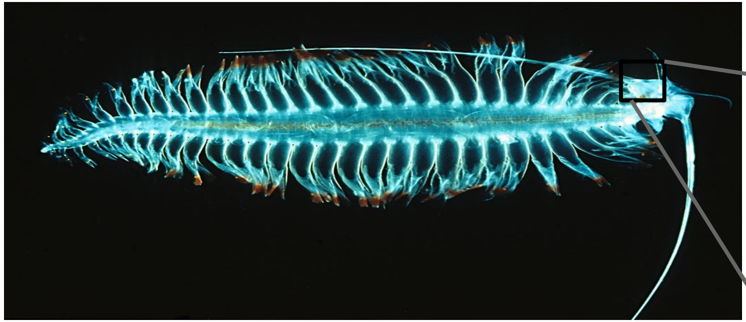

What we see

194	210	201	212	199	213	215	195	178	158	182	209
180	189	190	221	209	205	191	167	147	115	129	163
114	126	140	188	176	165	152	140	170	106	78	88
87	103	115	154	143	142	149	153	173	101	57	57
102	112	106	131	122	138	152	147	128	84	58	66
94	95	79	104	105	124	129	113	107	87	69	67
68	71	69	98	89	92	98	95	89	88	76	67
41	56	68	99	63	45	60	82	58	76	75	65
20	43	69	75	56	41	51	73	55	70	63	44
50	50	57	69	75	75	73	74	53	68	59	37
72	59	53	66	84	92	84	74	57	72	63	42
67	61	58	65	75	78	76	73	59	75	69	50

What the computer sees

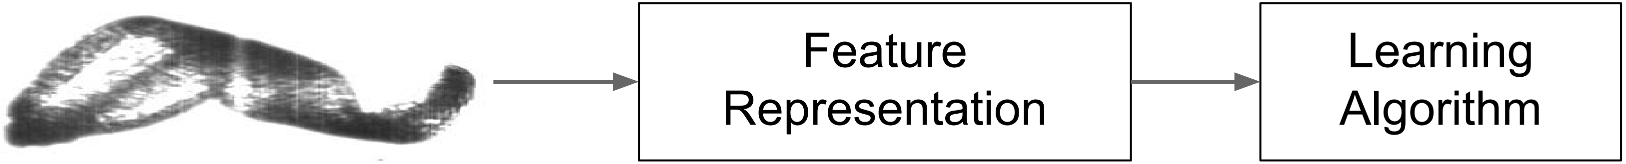

How to determine features given the image?

Features for vision

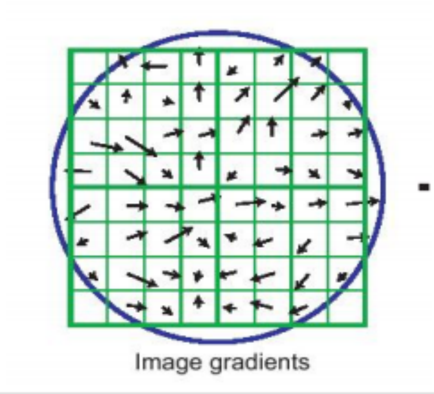

SIFT

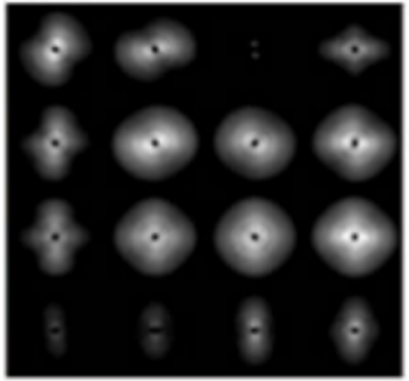

GIST

Domain specific hand engineered features like

- Ratio of glob's width and height
- Shape/Size

Learning the features!

Using Neural Networks (Inspired by nature)

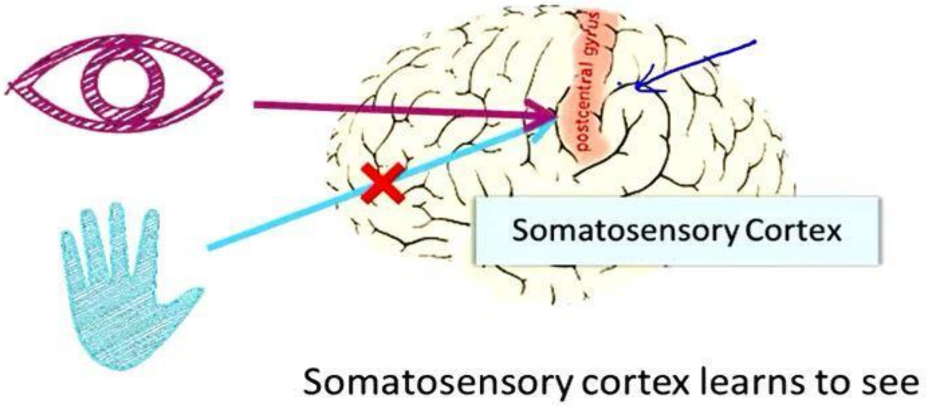

One Learning Algorithm Hypothesis

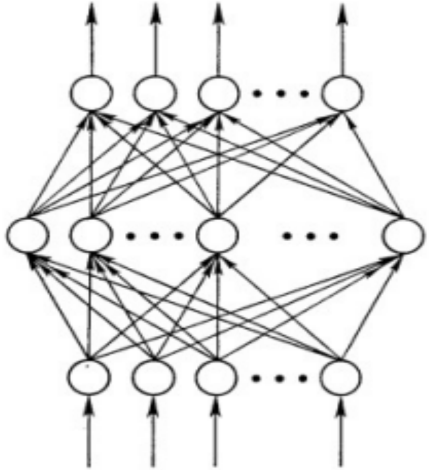

Neural Networks

Convolutional Neural Networks

Neural Networks with :

Local Connectivity

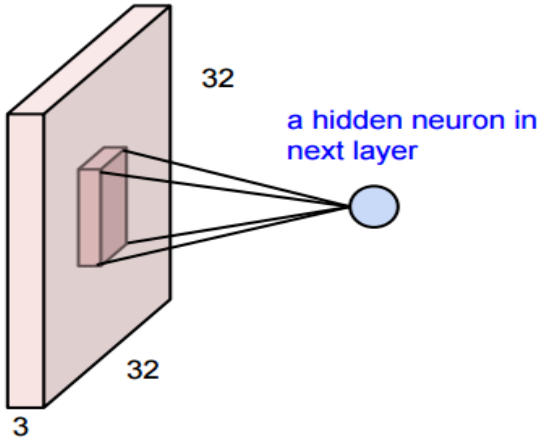

Same weight for neurons in a depth slice

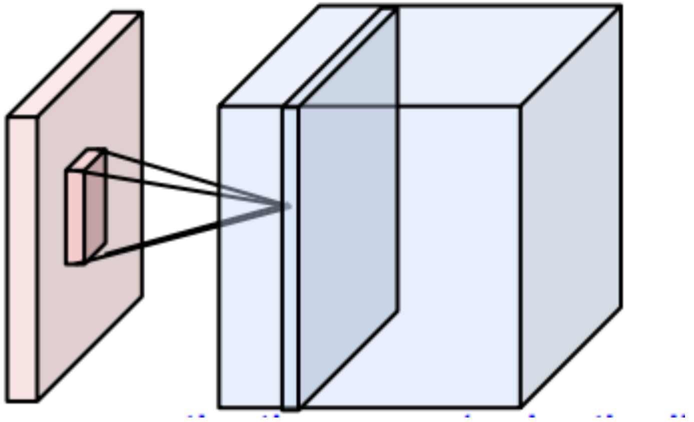

Layers used to build CNN

Convolutional Layer

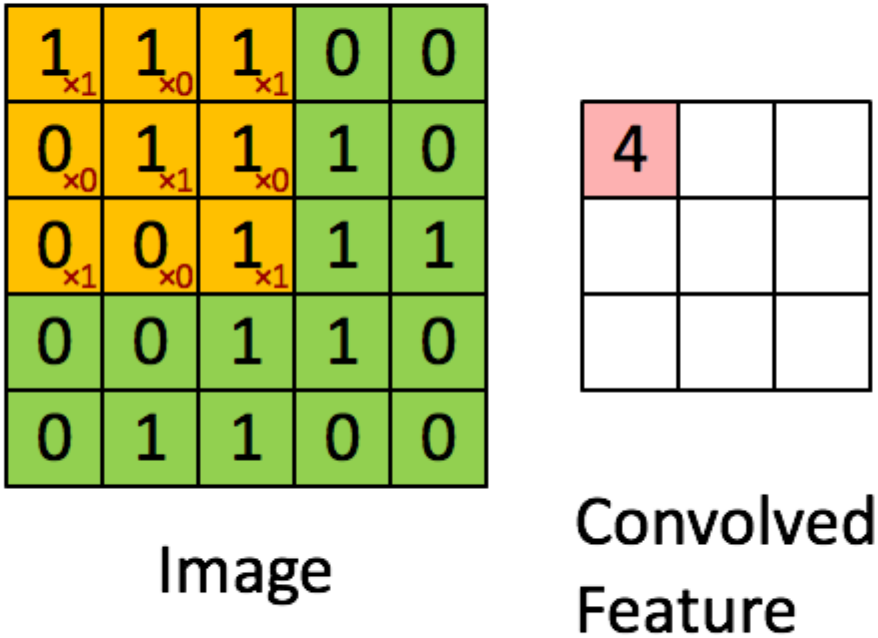

Polling Layer

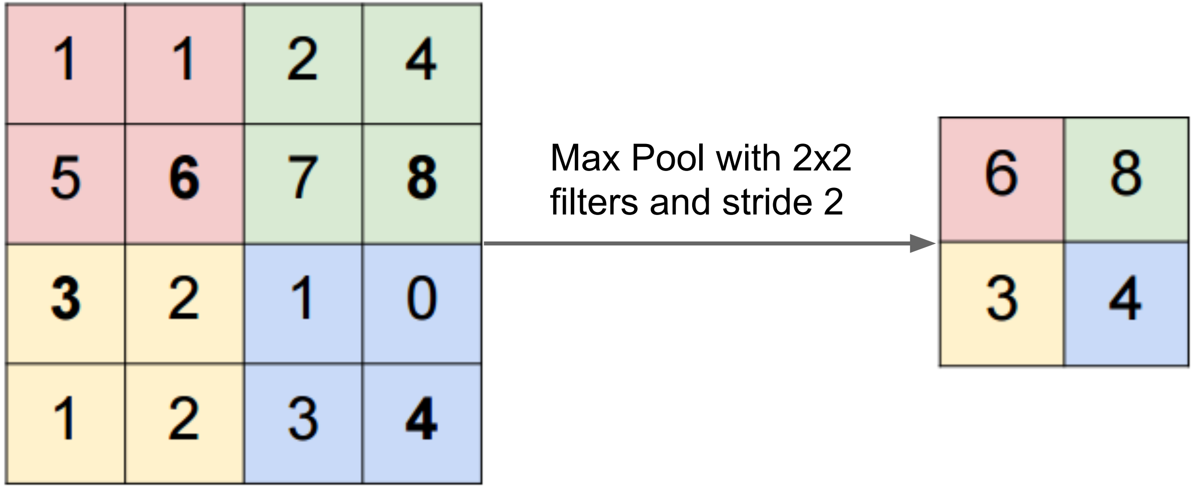

RELU Layer

Apply elementwise activation function such as max(0,x)

FC (i.e. Fully Connected) Layer

As with ordinary Neural Networks and as the name implies, each neuron in this layer will be connected to all the numbers in the previous volume.

CNN Example

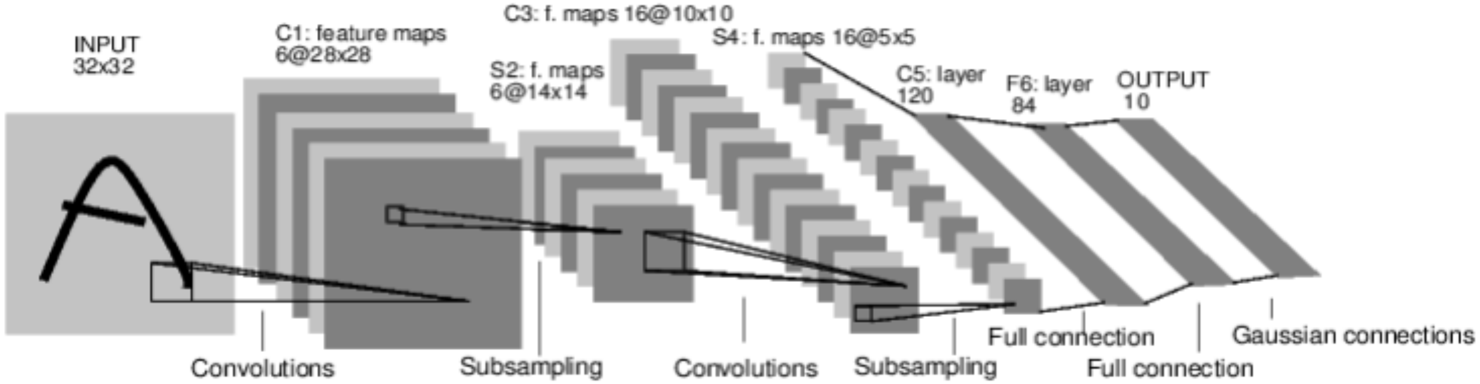

[LeNet-5, LeCun 1980]

Typical CNNs for vision look like

- [CONV-RELU-POOL]xN,[FC-RELU]xM, SOFTMAX
- [CONV-RELU-CONV-RELU-POOL]xN,[FC-RELU]xM,SOFTMAX

Work already done

- Explored the dataset
- Learnt to use AWS and used it to train a CNN
- Read some theory
- Tried Random Forest with hard coded features*

* Used the getting started code available online

Future Work

- Designing the Network
- Preventing Overfitting
  - Data Augmentation
  - Dropouts
- Benchmarking against SIFT

Why data augmentation?

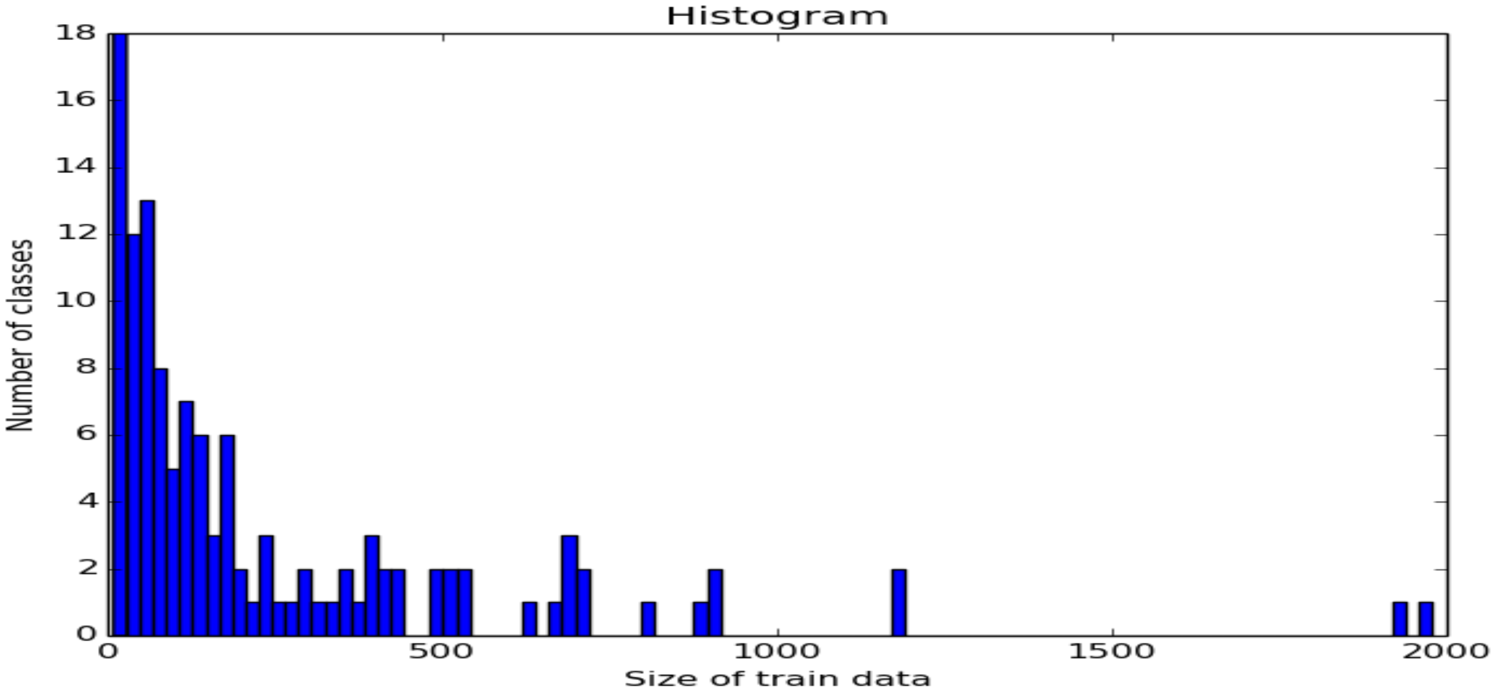

References

- Lecun Y. , Bottou L. , Bengio Y. , Haffner P. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11),2278 - 2324,1998
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems. 2012.
- Andrew Ng's Deep Learning Lectures

  http://cs229.stanford.edu/materials/CS229-DeepLearning.pdf
- CS231n : CNN for Visual Recognition Lectures

  http://vision.stanford.edu/teaching/cs231n/index.html

Questions?