

Fish Species Likelihood Prediction

Kumari Deepshikha (1611MC03)

Sequeira Ryan Thomas (1611CS13)

Content

1. Problem Statement
2. Literature Survey
3. Data Sources
4. Faster R-CNN – training
5. CNN – training
6. Results

Problem Definition

- Kaggle-The Nature Conservancy Fisheries Monitoring challenge.
- The objective is to **detect** and **classify** different species of fish which appears on a fishing boat, based on images captured from boat cameras of **various** angles.

Problem Definition

- Expected output is to predict the likelihood of fish species in each picture.
- **Eight target categories** are available in this dataset: Albacore tuna, Bigeye tuna, Yellowfin tuna, Mahi Mahi, Opah, Sharks, Other(meaning that there are fish present but not in the above categories),and No Fish (meaning that no fish is in the picture).

Literature Survey

1. Classification problem

CNN – Early Days

- Convolutional Neural Network have been shown to **eliminate** the need for **hand-crafted feature** extractors,
- Thereby reducing the need for hand-craft heuristics, manual labelling and manual parameter tuning.

CNN – Early Days





Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

[Source:](#) Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition

CNN – Early Days



[Source:](#) Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition

CNN – Used for complex problems

- With the amount of computation available and the advances in deep learning, even **complex problems** can be solved with relatively less effort.
- Irene Simie de Torres describes how CNN is used to detect deforestation in satellite images.
- The problem of identification of **deforested** areas in satellite images boils down to detecting patterns in the image, similar to our problem.

CNN – Used for complex problems


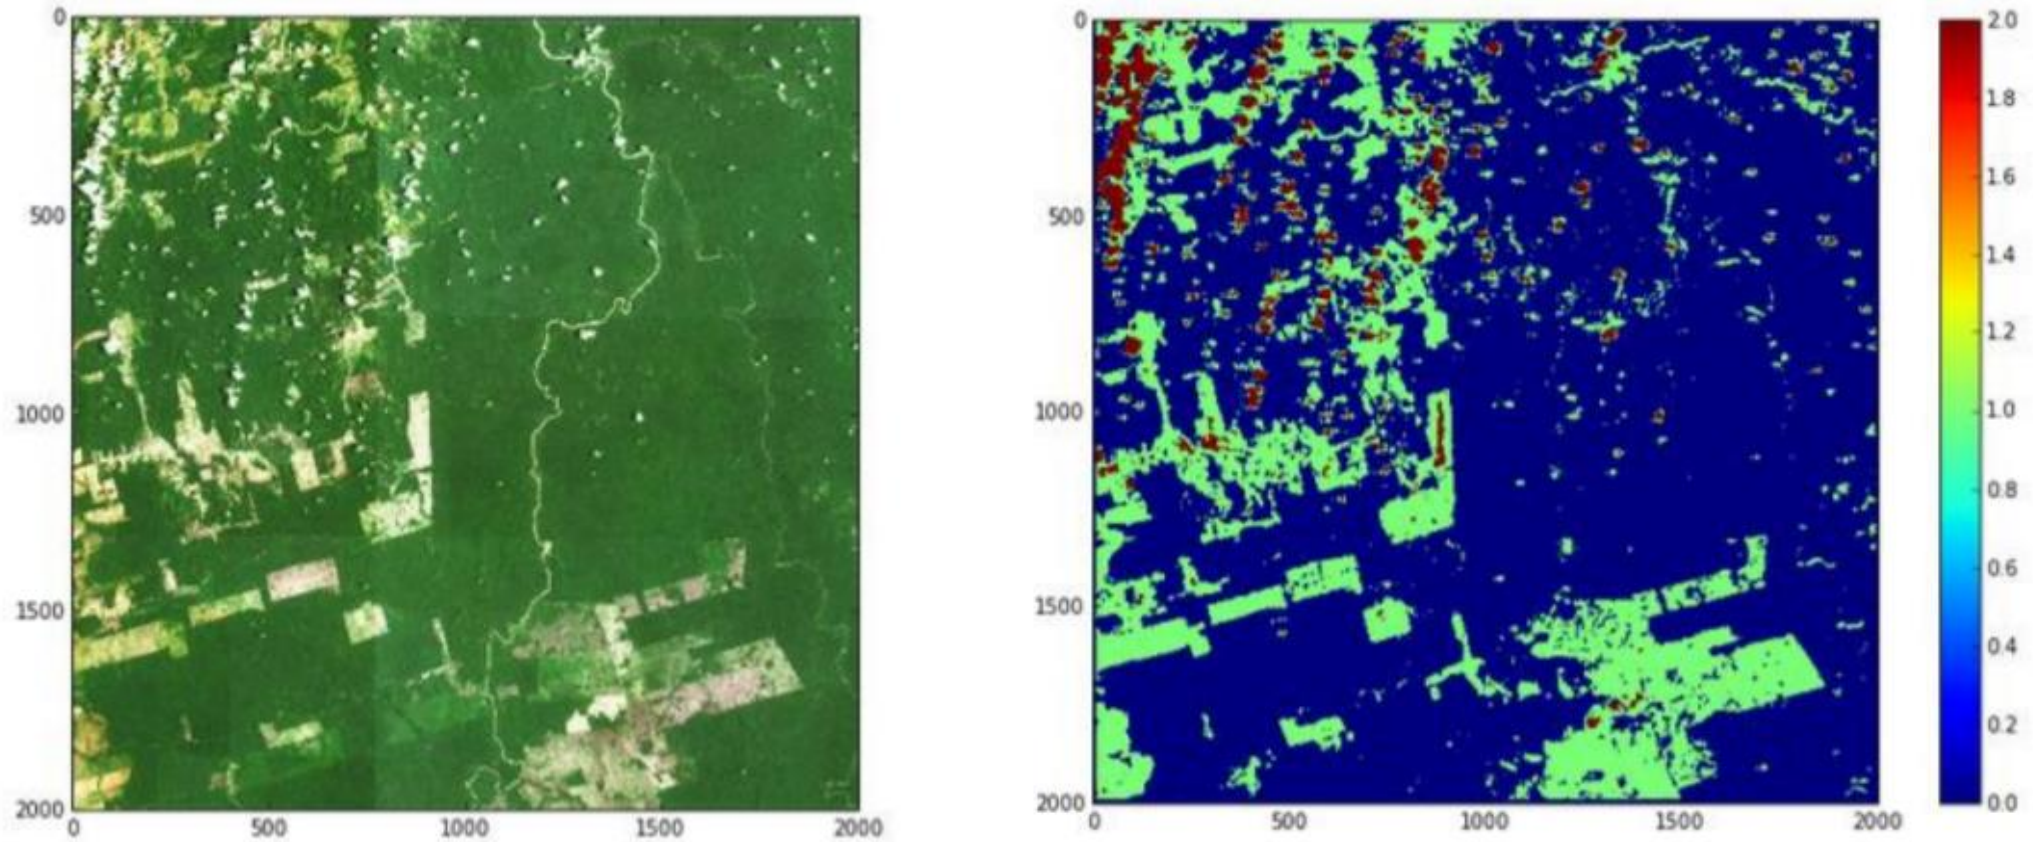


Figure 3.5: Representation of the model chosen

Source: Irene Simie de Torres. Analysis of satellite images to track deforestation

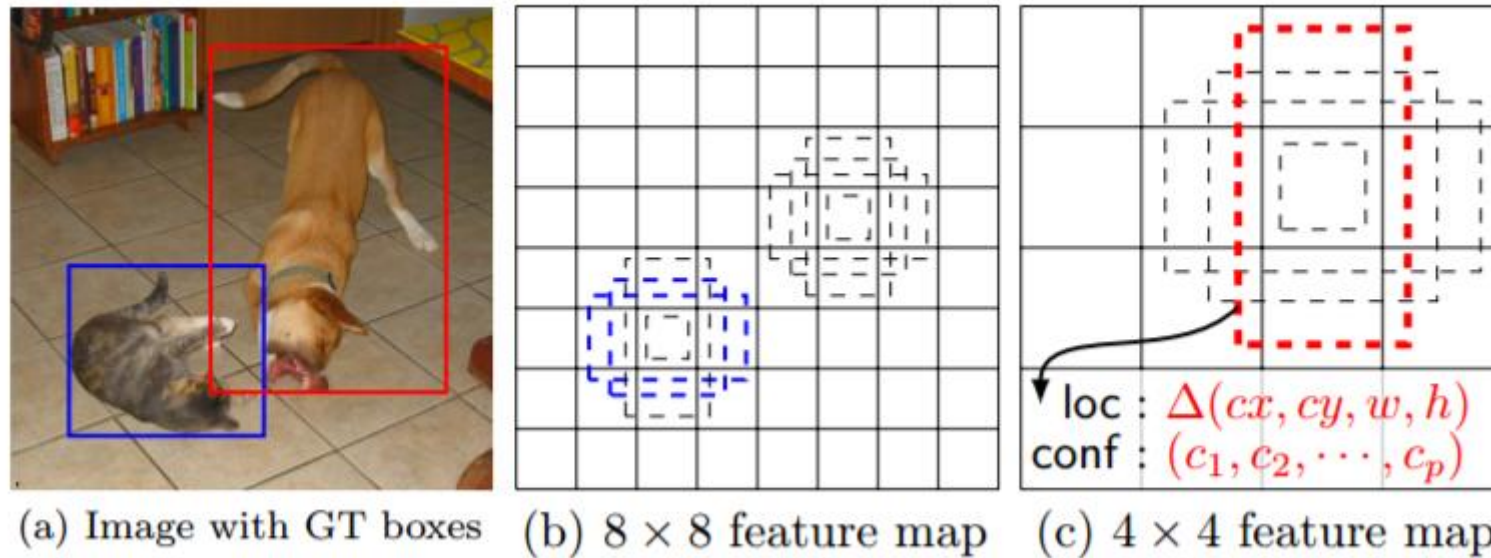
CNN – Used for complex problems



Source: Irene Simie de Torres. Analysis of satellite images to track deforestation

1. Localization problem

Single shot multibox detector (SSD)



Source: Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector.

Faster R-CNN

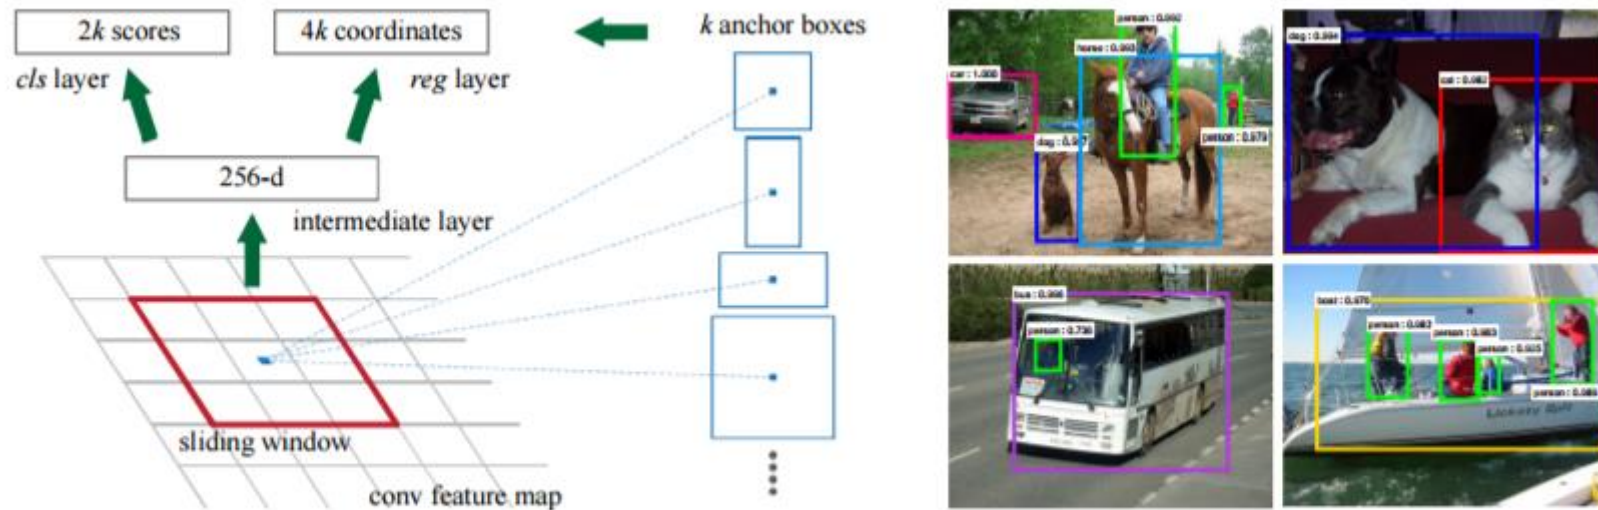


Figure 3: **Left:** Region Proposal Network (RPN). **Right:** Example detections using RPN proposals on PASCAL VOC 2007 test. Our method detects objects in a wide range of scales and aspect ratios.

Source: Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks

Localization problem results

- In both object detection models PASCAL VOC 2007 dataset was used for evaluation.
- Accuracy of 73.2% was achieved by Faster R-CNN and 76.8% was achieved by SSD model.

Coming back to our problem

Data Sources

- train.zip - zipped folder of all train images. The train folders are organized by fish species labels
- test_stg1.zip
- test_stg2.zip
- Size of data set - 2.1GB

Approach

- Used Faster R-CNN to detect regions of the image that contained fish.
- The cropped region was given to the CNN as input for species classification

Input to Faster R-CNN (Training)

data/train/LAG/img_04297.jpg,524,113,953,499,fish

data/train/LAG/img_04297.jpg,854,445,1134,711,fish



Input to CNN (Training)


- CNN was trained on cropped regions of the images containing fish.





Final output (testing)



Localizer Results



Classification Results



References

- [1] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278-2324, 1998.
- [2] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*, pages 21-37. Springer, 2016.
- [3] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91-99, 2015.

References

- [4] Kumari Deepshikha Sequeira Ryan Thomas. Git repository for fish species likelihood prediction. <https://github.com/seqryan/DeepLearningProject>, 2017.
- [5] Irene Simie de Torres. Analysis of satellite images to track deforestation. B.S. thesis, Universitat Politècnica de Catalunya, 2016.
- [6] Semmy Wellem Taju, Yu-Yen Ou, et al. Using deep learning with position specific scoring matrices to identify efflux proteins in membrane and transport proteins. In Bioinformatics and Bioengineering (BIBE), 2016 IEEE 16th International Conference on, pages 101-108. IEEE, 2016.

Thank You