



Hackathon Intel Image Classification

Adam, Brent, Zayn



TOC

Overview

Project objective

Market trends

Custom model

EfficientNetV2L

Issues with `image_dataset_from_directory`



Project objective

To create a convolutional neural network classifier that can differentiate between buildings, forests, oceans, mountains, streets, & glaciers

Overview

Data Collection

- 14,000 images in Train
- 3,000 images in Test*
- 7,000 images in Prediction**

Data Cleaning / EDA

- Dataset was clean, dealing only with images that were already sorted

Modeling and Interpretation

- Convolutional Neural Network
- Had most difficulty classifying Buildings against Streets, And Glaciers against Mountains
- 90.09% val_accuracy score





Image_dataset_from_directory is powerful...

Read in data directly from local folders

Automatically categorize images according to which folder they were in

Shift images to be the same size

Can also rotate, flip, and modify images to increase test dataset

Is the updated & combined version of Data Image Generator and from_directory

Rescaling of RGB values is now recommended to be done as the first layer of your CNN


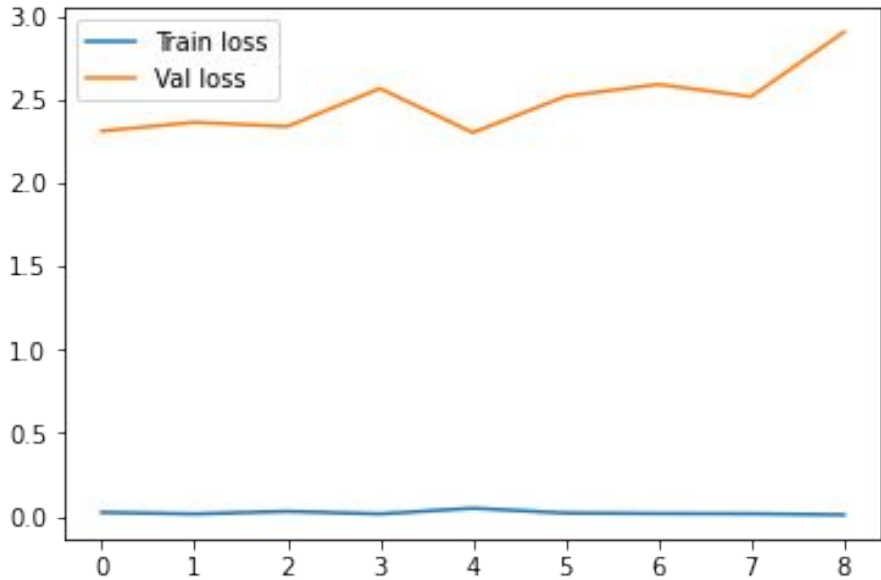
But there are some issues...



Issues with `image_dataset_from_directory`

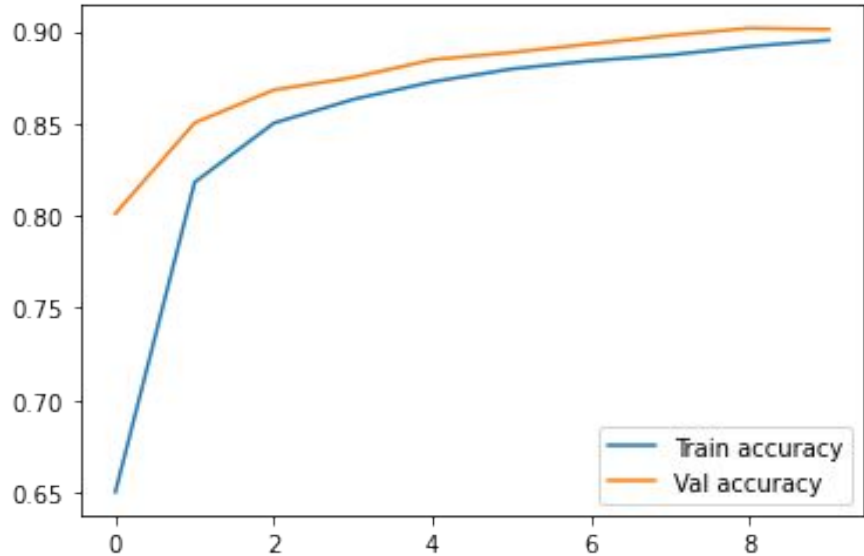

- We wanted to create a confusion matrix on unseen test data to validate which kinds of images the model was best and worst at predicting
- However, in creating that matrix, we noticed that Keras was only seeing 5 categories instead of 6 in the test data
- After extensive debugging, we couldn't solve the issue - so we got creative
- Implemented a train test split on the training data to avoid overfitting
- Created a confusion matrix on the entire dataset to be able to answer our original question - what kinds of images is our model best and worst at predicting?

Our custom model performed at 67% accuracy



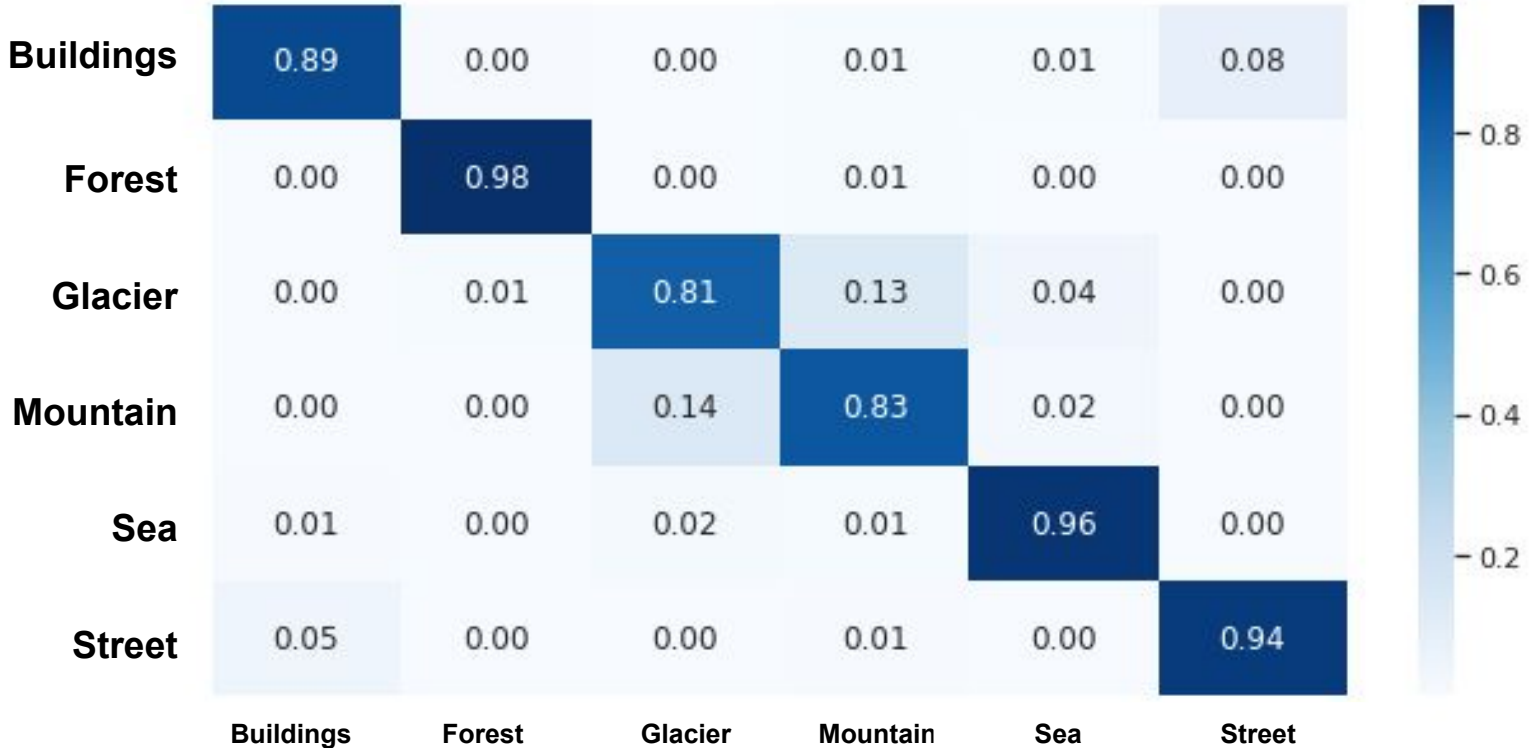
Baseline score for a 6 category classification problem is 16.6%

Pre-trained model achieved 90% accuracy



Baseline score for a 6 category classification problem is 16.6%

Model is best at predicting X, worst at Y



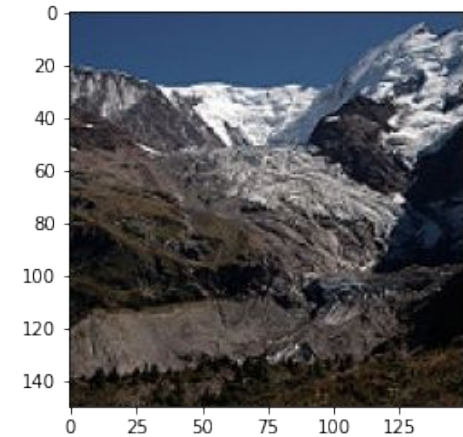
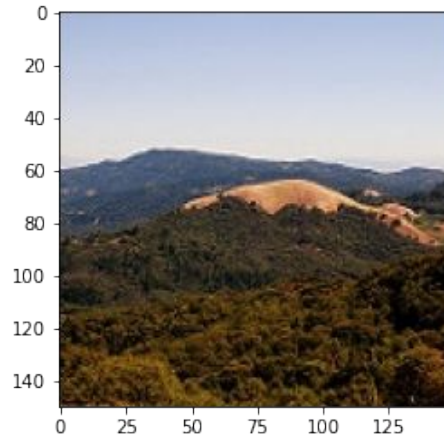


**To emphasize how
good 90% accuracy
is on this task, let's
play a game...**



Let's play a game: Which is which?

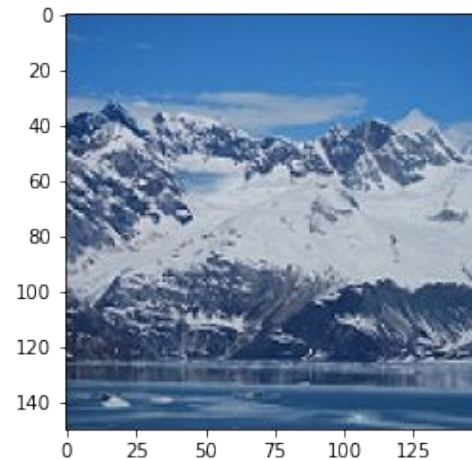
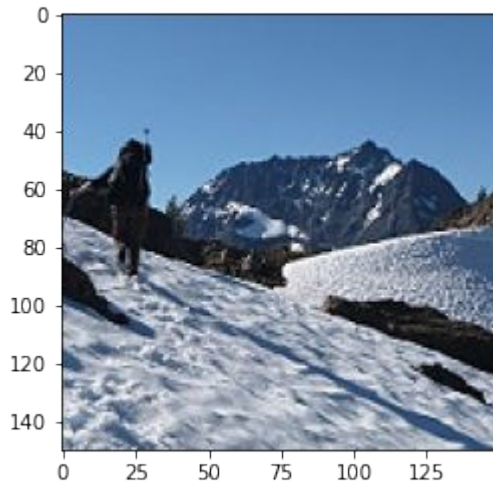
Mountain or glacier?





In these 2 pictures which is a glacier and mountain.

Can you spot the correct image or are humans not as superior to a simple Neural Network...



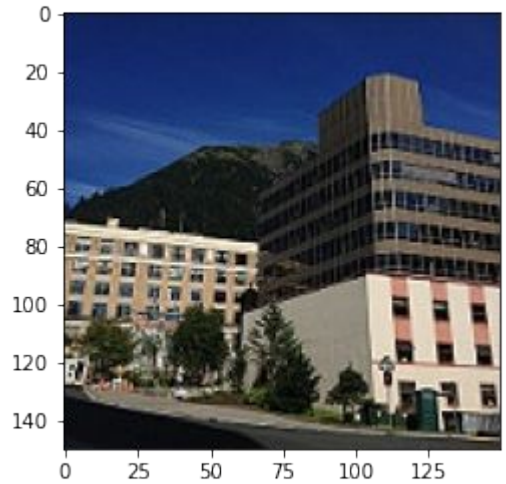
Let's play a game: Which is which?

Building or street?



Let's play a game:
Which is which?

Building or street?





Conclusion

Image classification is a powerful tool that requires complex architecture, extremely deep models with very large numbers of parameters, and huge datasets for effective training - but with the pre-trained models available in Keras, you can look like a pro with just a few tweaks and custom layers!





Thank you.

