

DHS/Kaggle Passenger Screening Algorithm Challenge

1st Place Solution

Jeremy Walthers

DataLab USA

jwalthers@datalabusa.com

Overview

Method

- Multi-View Convolutional Neural Network

Results

- 1st Place, 0.02417 logloss on test data, (for those not familiar with logloss this is roughly like 97.6% correct)

Motivation to Participate

- Learning, Experimentation, (prize pool impacted effort level)

Benefits Derived


- More familiarity with recent computer vision research, Practical experience implementing complex and distributed models in Tensorflow, (and \$500,000)

Future

- Would be interested in working on similar problems in the future (jwalters@datalabusa.com)

Technical Description

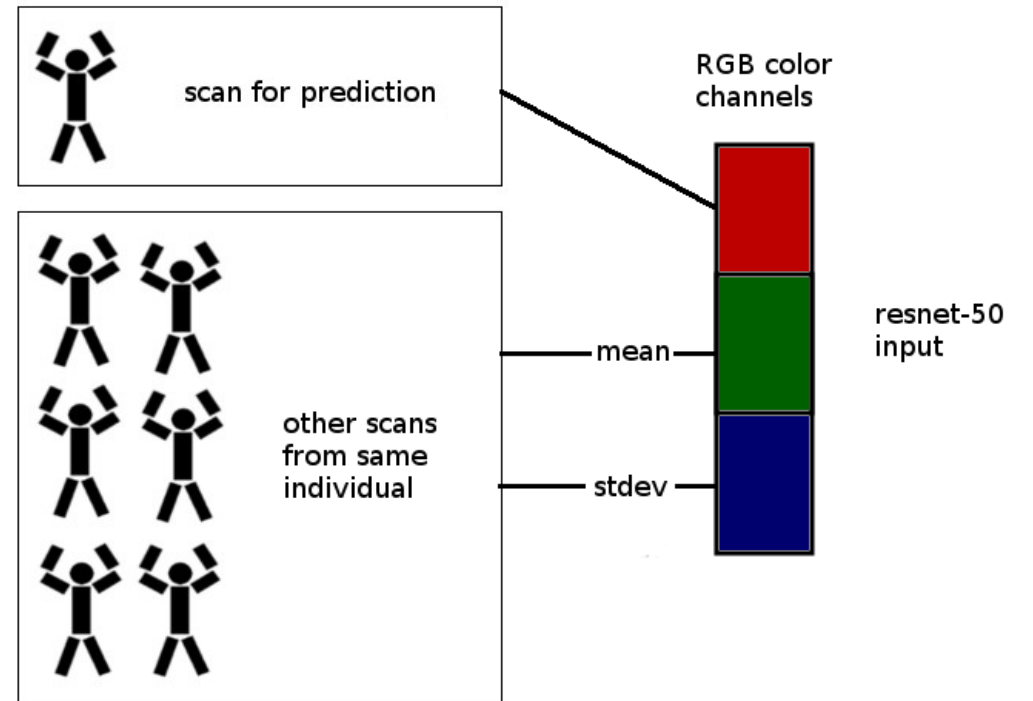
- Contest data provided in four formats
- Present solution used only APS format
 - APS = “projected image angle sequence”
 - Imagine walking in a circle around the subject and taking 16 photos
- APS was the smallest format at ~10MB per scan
 - Largest format, AHI, was over 2GB per scan
- Algorithm splits scans into four overlapping regions depicted in the figure to the right
 - Separate model built for each region
 - Allowed higher resolution models



Each scan split into four regions
separate model built for each region

Technical Description

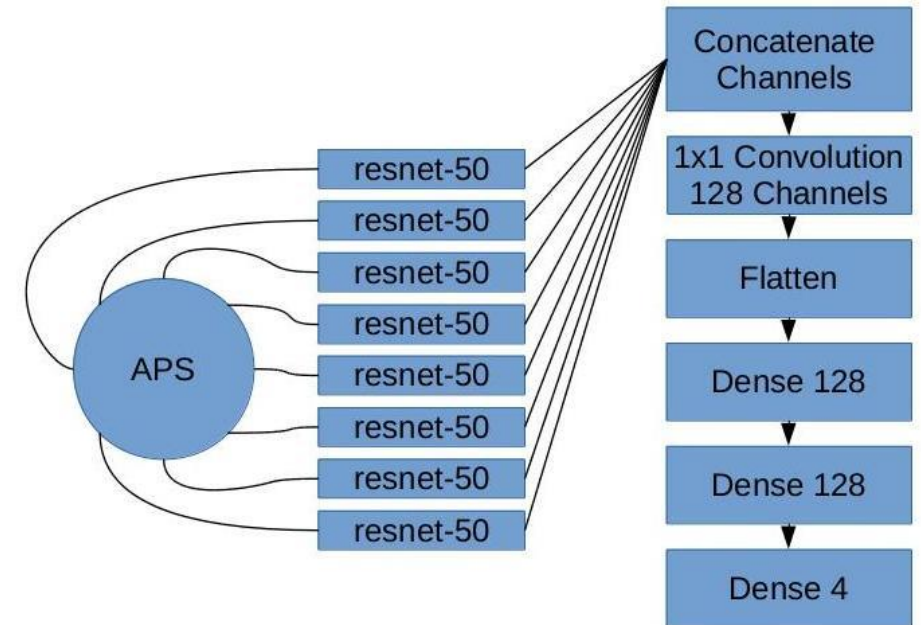
- Scans are monochrome
- Standard computer vision models take three-channel RGB input
- Contest data consisted of multiple scans of the same subjects
 - Rules permitted use of multiple scans in prediction
- Other scans for an individual identified via unsupervised learning
 - Mean and standard deviation are taken at the pixel level



Generating three-channel input from reflectivity














Technical Description

- Eight of Sixteen APS images fed into ResNet-50 and output reduced to four or five threat predictions on each region
 - Only half of the APS images used due to hardware constraints
 - ResNet-50 pre-trained on the ImageNet-1k dataset
 - Weights shared between each copy of ResNet-50
- Predictions pass through a calibration step before submission
 - Accomplished by what is typically called stacking
 - In this case: a gradient boosted decision tree model was fit on out-of-sample predictions of the model depicted to the right



Model architecture

Results

#	Δpub	Team Name	Kernel	Team Members	Score 	Entries	Last
1	▲ 134	idle_speculation		 ●●●●●	0.02417	2	10mo
2	▲ 71	serg14		 ●●●	0.02659	2	10mo
3	▲ 69	David O. Thomas A.		 ●●●●	0.03042	4	10mo
4	▲ 1	teedrz		 ●●●	0.04211	3	10mo
5	▲ 40	Oleg Trott		 ●●●	0.04236	2	10mo
6	▲ 116	CNN is fake model		 ●●●●●	0.05501	2	10mo
7	▲ 64	suchir		 ●●	0.05838	2	10mo
8	▲ 41	kaggle446		 ●●●	0.05970	10	10mo
9	▲ 108	Trox&Troy		 ●●●●●	0.06036	2	10mo
10	▲ 104	Moejoe		 ●●	0.06132	2	10mo
11	▲ 92	dhammack		 ●●●●	0.06438	4	10mo
12	▲ 125	Joseph Chui		 ●●●●	0.06783	2	10mo

Strengths and Weaknesses

Weakness

- Slow (~2 minutes per scan)
 - Can reduce to a couple seconds on GPU with very minor performance hit
- Method relied on multiple scans of the subject
 - Small (0.00-0.01) increase in logloss for single scan approach
- Variance in subject height caused prediction issues
 - Easy to fix

Strengths

- Low error rate
 - No errors when region is visible in stage 1 validation data
- Excellent generalization
 - Final standing mirrors cross-validated results
- Straightforward architecture

Future

Suggestions for future competitions

- More Data
 - Not in the sense of raw size, but in number of distinct training instances
- More Interaction
 - Take part in forum conversations
- Use Kaggle again
 - “Do it yourself” competitions frequently end unsuccessfully and unhappily

Further work

- Would compete in another DHS algorithm competition
- Interested to hear about other opportunities for similar work

Thank You

Jeremy Walthers

DataLab USA

jwalthers@datalabusa.com