

Security Coursework Feedback

Dr Chris G. Willcocks

Last Modified: May 3, 2019

Individual Feedback and Marks

Classifier Comment: This is a well-designed residual convolutional architecture, where the test accuracy was still significantly increasing; it is expected that leaving this running for longer would have resulted in a better test accuracy, its hard to give feedback on the architectures limitations without leaving it running for longer - maybe it could be a little higher capacity but the design looks good. The model itself has few parameters, but is deep and well regularised with dataset augmentation.

Generating a Pegasus Comment: The generative architecture is a typical convolutional autoencoder, where the sampling strategy adopted is to interpolate between a white horse and a bird with black wings. It is expected that a better result would be attained by sampling between a white horse and a white bird, (or horses with airplanes) both with similar background, or constraining the manifold such that it connects regions with similarly colored objects at the centre together (white horses with white birds). In terms of the dataset, some augmentation was adopted. It may have also been beneficial to throw away objects such as cars, ships, and trucks such that the interpolative space isn't interfered by such geometric regions of points. With better sampling or pre-processing, the model itself could have been made less blurry by increasing capacity and training for longer, or adversarial components could have been added and/or applied to shallow layers such as to capture higher frequency details such as textures and sharp edges.

Overall Comment: Overall this was a very good submission, where you had a particularly strong residual classifier which should just have been left to train for longer. The generative model architecture was acceptable, although more thought needs to be put into the sampling strategy; its unlikely to find a Pegasus interpolating between a white horse and a bird with black wings in 12 space; given all the other geometry of cars and trucks, its far more likely to just do some unmeaningful or linear interpolation between inputs without further constraining or connecting the manifold to more meaningful points (e.g. white horses with white birds, or white horses with white airplanes).

Student code:	gkgf37
Classifier ideas and experiments:	20/25
Classifier architecture:	10/10
Test accuracy: $\lceil (49\% - 20)/4 \rceil =$	8/10
Novelty:	4/5
Generative model architecture:	12/25
Recognisability:	6/10
Realism:	5/10
Uniqueness:	5/10
Overall mark:	70/100